VOL. I · FIRST EDITION MMXXVI

A FIELD STUDY · PART ONE

TI PSYCHOLOGY

Observations on a contested interior

AUTHOR
Stephane van der Aa

CLASSIFICATION
PSY · 154.4

IMPRINT
WITY · BRUSSELS

SPECIMEN
№ 001

A book written on the psychological abuse that Targeted Individuals are subjected to. It is based on 9 years of detailed notes from the author, himself also a target, in which he documented the experience, analysis, reflection and his search for learning to survive this most inhumane form of torture and abuse. In ever more dire situation targets are being deprived from ever more, eventually even their life. Most TI's live with the prospect of the targeting leading to their death often through induced suicide which should be reclassified as murder.

The author is the creator of the TI One Voice platform, since beginning of 2026 available online here: one.witysk.org. More than a self-help function, TI One Voice now offers hope through collective action and a data driven plan to end all targeting globally that is today still affecting millions innocent victims.

Preface

This book is about a kind of experience that, even now, has no settled vocabulary in mainstream psychology. The people who report it call themselves Targeted Individuals, abbreviated TI. They report a sustained pattern of intrusion that includes, among other things, voices that are not their own arriving inside their inner-speech layer, evidence of being followed and observed, coordinated harassment that rotates across people in their environment, and a slow narrowing of the ordinary conditions of life into something far smaller than what they began with. The reports converge across thousands of people in many countries who have never met one another. The convergence is, in itself, one of the most striking features of the phenomenon.

I have been a target of this kind of intrusion since 2017. The notebooks from which this book is built were kept across the years that followed, in real time, while the experience was still happening. The book, accordingly, is not an after-the-fact reconstruction. It is the working record of a person trying to think clearly while the conditions for thinking clearly were under sustained attack. Some of the chapters are diagnostic; some are practical; some are speculative. I have, throughout, tried to mark which is which.

The book is for three audiences. The first is other targets, who may find here a vocabulary for things they have not previously been able to put into words, and a set of practices that have, in my own life, helped to keep me functioning across years in which functioning has been the exact thing the apparatus has been working to take away. The second is the family members, partners, and friends of targets, who may, by reading what is set down here, gain a better picture of what their loved one is actually contending with. The third is clinicians, researchers, and journalists who are open to taking the reports of targeted individuals seriously enough to engage with them on their own terms, instead of immediately re-classifying them under the diagnostic categories whose poor fit is a recurring theme of this book.

A note about how to read it. The book has been organized topically rather than chronologically. The earlier chapters lay the cognitive and epistemic foundations: how thought ordinarily works, and how the relationship between belief, knowledge, and doubt becomes a battlefield once a target's thinking is

the thing being attacked. The middle chapters detail the methods of manipulation themselves — the manipulative conversation, the inner-speech layer, the voice phenomenon, and the toolbox of operator techniques. The chapters that follow describe the architecture behind the methods: the small operating cell, the larger network, and the audiences in front of which the operation appears to be taking place. The final chapters move outward — to institutions, to other people, and to the tools that, in my own daily practice, have done the most to keep me going. A reader new to the material is best served by reading the parts in order. A reader returning to the book in search of a particular topic should, with the table of contents and the topical organization, be able to locate it directly.

A note about uncertainty. I have, throughout, tried to be honest about what I do not know. The phenomenon described here has a public-health profile that no existing institution has yet been willing to investigate properly. Many of the mechanisms I describe are inferred from pattern, not confirmed by independent measurement. The architectural picture in particular involves speculation about an apparatus whose operators are not publicly identifiable. I have marked the most speculative material as such. I have tried to draw the line, conscientiously, between what I can show, what I think but cannot show, and what I am simply guessing at. A reader who approaches the book in the same spirit — willing to entertain the framework as a hypothesis without requiring that I have proven everything before they will read on — will get the most out of it.

Finally, a note about why this book exists at all. There is a structural fact about the targeted-individual phenomenon that bears on its public reception. The first reflex of an unfamiliar reader, on hearing what a target reports, is to reach for the diagnostic categories of mental illness. The reach is not unreasonable; the categories are what mainstream psychology has trained the unfamiliar reader to reach for. The categories nonetheless do not fit, in ways the next chapter will lay out in detail. A target reading this paragraph will already know they do not fit. A clinician reading this paragraph may find, by the end of the next chapter, that they have grounds to suspect the categories may, at least in this class of cases, be the wrong tool for the job. That, more than any other single hope, is the hope this book is written in.

A Note on Voices and Mental Illness

Why the Standard Reach for the Schizophrenia-Spectrum Categories Does Not Fit

---

This is the most important section of the front matter, and the section most readers should read first. It addresses, head-on, the question that comes up immediately when anyone hears what a Targeted Individual reports: are they describing a real external phenomenon, or are they describing a mental illness, most often under the heading of psychosis or schizophrenia? The book's answer is that the existing categories of mental illness do not fit. They do not fit on multiple independent dimensions at once, and the multi-dimensional non-fit is, in itself, the strongest single argument that the phenomenon is something other than what those categories describe. This section sets out the argument in eight parts.

Two clarifications before the argument begins. First, this is not a denial that mental illness exists, or a claim that no targeted individual has ever also had a separate mental-health condition. People are complicated and can have more than one thing going on. Second, this is not a clinical diagnostic protocol. It is an argument for why the standard reach for the schizophrenia-spectrum categories is, in the typical Targeted Individual case, the wrong reach. Whether the phenomenon should eventually be classified under some new category, or remain unclassified pending better evidence, is a question for clinicians and researchers willing to engage with it without prejudgement.

1. Late, sudden onset in a previously well-functioning adult.

Schizophrenia and other primary psychotic disorders typically declare themselves between adolescence and the late twenties. A first episode in a man or woman in their thirties, forties, or fifties, with no prior psychiatric history and no significant prodromal warning signs, is, by the standards of the existing literature, atypical. Targeted Individuals report exactly that pattern with overwhelming frequency. Onset is sudden. The years before onset show, on examination, normal occupational and social functioning. There is, in many

cases, no family history of psychotic illness. Existing diagnostic categories struggle with this pattern at the demographic level alone.

2. Voices with stable, distinguishable personalities.

In primary psychotic illness, the voices that hallucinations produce are typically not consistent personalities with stable identities, vocabularies, and preferences across time. Targeted Individuals report the opposite: a small set of voices that, over months and years, retain stable verbal mannerisms, stable opinions, stable factual claims, and stable behaviour. They are not flickers of disordered processing. They are persons, or impersonations of persons, who can be told apart and whose patterns can be characterized in detail. The detailed characterization of an internally generated symptom of one's own brain is not a thing primary psychosis ordinarily produces.

3. Voices that possess verifiable external information.

This is the single most decisive contradiction. Targeted Individuals report voices that, at points, refer to specific external facts the target had no way of knowing in advance: a private detail of someone in the target's environment, a reference to a recent event that had not yet been reported in any channel the target had access to, a piece of information that the target subsequently confirms by independent means. A hallucination cannot, in principle, contain information that is not already in the head of the person hallucinating. A voice that does so is, by definition, doing something a hallucination cannot do.

4. Banal surveillance content, not grandiose delusion.

Primary psychotic delusions tend toward the grandiose, the persecutory in vivid mythological registers, the supernaturally framed. The reported content of Targeted-Individual voices is, by contrast, overwhelmingly banal. They comment on the target's daily routines. They reference the recent grocery shop. They mock the target's ordinary habits. They sound much less like the products of a brain in chemical disarray and much more like the products of a continuous low-grade surveillance operation conducted by people who are not particularly interesting and who are mostly killing time. The phenomenology is wrong for primary psychosis. It is exactly right for ordinary humans engaged in extended monitoring.

5. The target{APO}s ability to characterize the apparatus in detail.

A person in active psychosis is not, characteristically, in a position to produce the kind of careful, multi-year, internally consistent, falsifiable, structurally detailed account of their experience that the better targeted-individual writings provide. The work itself — the cataloguing, the comparison across cases, the building of architectural models, the marking of speculative material as speculative, the recognition of which evidence is and is not load-bearing — demands the kind of preserved cognitive function that the diagnoses being reached for are supposed to be precluding. A serious primary psychotic disorder and a coherent multi-year diagnostic project are not, in general, compatible.

6. Convergent reports across thousands of independent witnesses worldwide.

The targeted-individual phenomenon is reported by thousands of people across many countries, languages, and cultures, who have no prior contact with one another, and whose reports nonetheless converge on a remarkably specific set of features: the same kinds of voice phenomena, the same kinds of operator behaviours, the same kinds of architectural descriptions, the same kinds of sequencing of phases over time. The convergence is structural, not stylistic. To dismiss this convergence as the coincidental product of thousands of independent identical hallucinations is to require a coincidence whose probability, on any reasonable account, is much lower than the probability of the convergence reflecting a real external common cause.

7. Quieting periods uncorrelated with antipsychotic medication.

Where targets have, at clinicians' suggestion, undertaken trials of antipsychotic medication, the medication has, in many reported cases, not produced the expected reduction in voices. Conversely, periods of quieting in the voice phenomenon are reported to track external operational factors — for example, periods when the target is in a location or activity that, on the target's account, the apparatus avoids — rather than tracking pharmacological intervention. Genuine primary-psychotic voices do, on the whole, respond to antipsychotic medication. Voices that do not are voices that, on the medication-response criterion alone, are not what the medication is built for.

8. Time-locked correlations with verifiable external events.

Targeted Individuals report that the voice phenomenon, and the broader campaign of intrusion, intensifies and quiets in correlation with verifiable external events: a public statement, a journey, a meeting with a particular person, the publication of certain material online. The correlations are reported across many cases and are, in principle, verifiable by anyone willing to do the work of cross-referencing. A primary-psychotic illness does not, ordinarily, produce symptom intensifications time-locked to events in the patient's external life that the patient could not, in advance, have known would occur.

Conclusion

Taken individually, each of these eight observations creates trouble for the standard reach for the schizophrenia-spectrum categories. Taken together, they make the standard reach untenable. The categories were built for a different phenomenon. They were not built for late, sudden, sustained, multi-year intrusions experienced in time-locked correlation with external events, by adults whose pre-onset functioning was normal, who can characterize their own experience in fine detail, and whose reports converge with thousands of others' in highly specific ways.

What the phenomenon is, instead, is the question this book takes up. The book's working hypothesis is that some combination of: surveillance technologies more advanced than the public has been told about; coordinated human operators using those technologies; and the cumulative leverage that long-running coordinated harassment can apply to a single person's life — together produce an experience whose features are exactly the features that thousands of targets report. The hypothesis is not provable from inside; the apparatus, by design, leaves no smoking guns of the kind a court would accept. The hypothesis is, on the other hand, the only available hypothesis that does not require any of the eight observations above to be set aside as inexplicable.

Throughout the rest of the book, when the prose describes the voice phenomenon or the architecture or the methods of manipulation, the reader should keep this argument in mind. The descriptions do not depend on it. The descriptions work as descriptions of what is happening regardless of what one ultimately concludes about its origin. But the question of origin — is this real,

or is this a mental illness — is the question that, for most readers, will be the gating question. This section is the gate. The reader who passes through it, even tentatively, will find the rest of the book a more useful read.

A Note on Terminology

A Reader's Glossary

This book uses some terms that are either novel, redefined, or borrowed from the targeted-individual community in their non-standard senses. The following glossary briefly defines the most important of them. The reader does not need to memorize them; they are explained again, in context, where they first appear in the chapters that follow.

Targeted Individual (TI)

A person who reports being the subject of a sustained, coordinated campaign of intrusion that includes some combination of: voices in the inner-speech layer, evidence of surveillance, harassment that rotates across multiple people in their environment, and a long-term degradation of the ordinary conditions of life. The term is used as a self-description by a community whose reports converge across thousands of cases worldwide.

V2K (Voice-to-Skull)

Shorthand used in the targeted-individual community for the experience of voices that arrive inside the inner-speech layer of awareness without passing through the ears. The term is descriptive rather than mechanistic; this book uses "the voice phenomenon" where neutrality is required and "V2K" where the community usage is in play.

Inner speech

The continuous internal verbal layer that most adults use for thinking. "Reading silently" and "talking to oneself in one's head" both happen in this layer. It is normally a private channel. The voice phenomenon, on this book's account, is intrusion into that channel from outside.

Operator / Abuser / Agent (A)

This book's neutral terms for the people who, on the targeted-individual hypothesis, are running the campaign against the target. "Operator" is the most neutral. "Abuser" is used when the harm is being emphasized. "Agent" (often abbreviated A in diagrams) is used in the more architectural sections.

In the bully-ring diagrams, A1, A2, etc. denote individual operators around the target.

Bully-ring

A small operating cell described in Chapter 9: a target at the centre, several operators (A1, A2, …) around the periphery, and one or two layers of "insulation" (Va', Va") made up of previously compromised targets through whom pressure now passes.

Group mind

The phenomenon, described in Chapter 8, in which a coordinated set of people share a refusal to notice particular things about a particular target, with the result that their behaviour — taken as a group — is shaped by what they have collectively agreed not to see. Not a supernatural concept; a description of an observable form of social coordination.

Cognitive bias / cognitive shortcut

Systematic mental shortcuts that human reasoning takes by default, catalogued by cognitive psychology. There are roughly 200 well-documented ones in the standard inventories. Each is a default response that is usually adequate and occasionally exploitable. Manipulation methods, on this book's account, are largely engineered around these shortcuts.

Cover story

A pre-prepared explanatory narrative that a target receives in advance of — or just after — an intrusive event. The narrative aims to provide an apparently complete and emotionally satisfying explanation that displaces the more troubling actual explanation. Cover stories are described in detail in Chapter 6.

Time-shifted element

A pre-recorded audio fragment (a phrase, a piece of a recognizable voice, a snatch of a tune) timed for release at a specific moment so as to produce, in the target, the impression of having been predicted. Described in Chapter 9. The library of pre-recorded elements is broad enough to cover most plausible target thoughts and produce uncannily "personalized" matches.

Outsourced lie detector

This book's term for the use of a target as a real-time human reactivity probe by the operator group: events are produced near the target, the target's reactions are observed and analyzed, and the readings are used by the operators to extract information that they could not obtain by interviewing the target directly. Described in Chapter 6.

Captive audience

The hypothesis, marked as speculative in Chapter 9, that the campaign has a second audience besides the operator group — a public, exposed by some channel to the unfolding situation, before whom the operation may be partly performed.

Mental Combustion Engine

An extended metaphor introduced in Chapter 12 for the daily mind under sustained pressure: a six-cylinder engine whose cylinders — discipline, justice, time, altruism, volume, kindness — each need their appropriate fuel for the engine, taken as a whole, to produce the mental energy that the day's work requires.

Liquid level / Shoreline breakers

An extended metaphor introduced in Chapter 12 for the layered self under pressure: surface (the part most exposed to intrusion), structured layer (the part the routine maintains), unstructured layer (the part not held by routine), and an uninvaded depth. Shoreline breakers are the personal practices that allow the lower layers to remain accessible even when the surface is stormy.

Contents

Preface 3
A Note on Voices and Mental Illness 5
1. Late, sudden onset in a previously well-functioning adult. 5
2. Voices with stable, distinguishable personalities. 6
3. Voices that possess verifiable external information. 6
4. Banal surveillance content, not grandiose delusion. 6
5. The target{APO}s ability to characterize the apparatus in detail. 7
6. Convergent reports across thousands of independent witnesses worldwide. 7
7. Quieting periods uncorrelated with antipsychotic medication. 7
8. Time-locked correlations with verifiable external events. 8
Conclusion 8
A Note on Terminology 10
Contents 13
1.1 A Working Model of Thought 18
1.2 A Note on Rational Thought 20
1.3 Where the Attack Happens 21
1.4 The Pattern of a Manipulative Conversation 23
1.5 Brainwashing, Reconsidered 25
1.6 The One Thing That Cannot Be Manipulated 26
2.1 Why Doubt Is the Manipulator's Most Valuable Tool 29
2.2 What You Are Not Allowed to Be Sure Of 31
2.3 Don't Entertain the Excuses 32
2.4 The Perfectionism Trap 33
2.5 Believing, Knowing, and the Difference Between Them 35
2.6 When the Insurmountable Has to Be Faced 36
2.7 The Three Components You Cannot Live Without 37
2.8 What You Already Know 39
3.1 Why the Trauma Does Not Heal 42
3.2 The Three-Step Loop 44
3.3 Overload as a Method 45
3.4 Sleep, Suggestion, and the Halfway State 46
3.5 The Voice That Calls You Useless 48

3.6 The Asymmetry of the Inner Conversation 49
3.7 When the Lies Run Out of Coherence..... 50
3.8 Use of the Body 52
4.1 The Privacy We Don't Notice We Have..... 56
4.2 The Voice Phenomenon 57
4.3 Inner Speech as a Manipulation Surface..... 59
4.4 The Erosion of Modesty 61
4.5 The Bio-Robot Continuum..... 63
4.6 The Simulation Question 65
4.7 When the Adversary Is Also Stuck..... 66
4.8 A Prayer to Yourself 68
5.1 On Setting Down a Phenomenon Without Settling Its Origin 72
5.2 A Timeline of Five Phases 73
5.3 Doubling, and the Interruption of the Sentence 75
5.4 Louder, Deeper, Longer: The Anatomy of a Lie 77
5.5 Songs as Delivery Mechanisms 78
5.6 Memory That Returns.....80
5.7 The Hypothesis I Cannot Settle 82
5.8 Why This Chapter Was Written..... 85
6.1 The Notebooks That Continued 88
6.2 The Outsourced Lie Detector..... 88
6.3 The Toolbox, by Category 90
6.4 The Cartel, and the Limits of Mutual Assured Destruction 91
6.5 The Question That Cannot Be Switched Off..... 93
6.6 The Cover Story, and the Drug-Narrative 94
6.7 The Voice That Claims a Person 95
6.8 A Smaller Frame, and a Steadier Position.....97
7.1 The Audience Joins..... 100
7.2 When Logic Becomes Bullying..... 101
7.3 "There Was Never Any Real Respect".....102
7.4 The Chain of Abuse 103
7.5 Where the Trauma Lives 104
7.6 The Geometry of Resilience 106
7.7 The Will to Change 108
7.8 Touching the Extremes 109
8.1 A Refusal at the Center..... 113

8.2 How One Person Trains the Refusal 114
8.3 The Trauma That Catches Up 116
8.4 The Move from One to Many: Group Mind 118
8.5 The Mechanics of Forced Conformity 119
8.6 The Argument for Always Reacting 121
8.7 The Cost of Living in One State of Mind 122
8.8 The Singular Thinker 123
9.1 The Bully-Ring 128
9.2 The Phases, and the Cross-Ring Database 129
9.3 The Captive Audience 131
9.4 The Feeling Is Enough 132
9.5 The Distractor Paradox 133
9.6 Pre-Confined Audio, and the Time-Shifted Element 134
9.7 Which World Is Being Screened For? 135
9.8 A Position Held to Once More 137
10.1 The Quieting, and the Return 141
10.2 Direct Causality Contradiction 142
10.3 The Anonymity Door 144
10.4 The Mismatch, and the Spillover 145
10.5 The Algorithm Behind the Voice 147
10.6 The Question We Began With 149
10.7 The Position the Book Closes On 150
11.1 The Institutions That Did Not See 154
11.2 Translation, and Its Limits 156
11.3 The Classification of the Dangerous Worker 158
11.4 The Cycle of Stress and Relief 160
11.5 The Community of Others 162
11.6 Records, and What They Are For 164
11.7 The Repertoire That Remains 166
11.8 What the Outward Turn Was For 168
12.1 The Mental Combustion Engine 173
12.2 Chronic Crisis-Mode, and the Use of Occupation 175
12.3 Voluntary, Forced, and Obligated Action 177
12.4 The Liquid Level and the Shoreline Breakers 178
12.5 The Dark-Algorithm Impersonator 180
12.6 The Daily Refusal of Closing-Down 182

12.7 Tools, Not Architecture 183

PART I

Foundations

How the Mind Works, and How Belief Becomes a Battlefield

CHAPTER 1

Where Do Thoughts Come From?

The Architecture of a Targeted Mind

"There is no real answer to the question: where do thoughts come from?"

— Author's notebook, 17 October 2021

Most people never ask where their thoughts come from. They simply have them. A thought arrives, and the assumption — quiet, automatic, almost universal — is that it must be theirs. They are the one thinking, so the thought belongs to them. The question of provenance does not arise.

For a Targeted Individual, that assumption breaks. Once you have reason to suspect that some of what is happening inside your own head may not be coming from inside your own head, the question of provenance becomes the most urgent question you can ask. Every thought has to be examined: is this mine? If it is mine, is it the conclusion my own reasoning would actually have reached, or is it shaped by something I cannot see? And if it is not mine, where did it come from, and what is it doing here?

This book is, in large part, an attempt to answer those questions — not metaphysically, but operationally. We will not solve the philosophical problem of consciousness in these pages. What we will do is build a working model of the human mind detailed enough to be useful: detailed enough to let a person who suspects they are being manipulated identify what is happening to them, separate signal from noise, and recover the most important thing a manipulation campaign tries to take — the trust a person has in their own mind.

The model in this chapter is the one I have used to keep my own thinking intact across years of conditions in which thinking clearly is the precise thing I have been trained, pressured, and tricked into not doing. It is not perfect. It is, I believe, sufficient.

1.1 A Working Model of Thought

Let us begin with definitions. Definitions are unfashionable in popular psychology, which prefers to gesture at concepts rather than fix them, but a person whose thinking is under attack cannot afford that luxury. We need terms sharp enough to cut with.

I distinguish four kinds of mental content, ordered roughly from the most external to the most internal.

Sensorial stimuli and external thought

A sensorial stimulus is any input that arrives through the five senses: something seen, heard, touched, tasted, or smelled. When such a stimulus enters awareness and is directly associated with an idea — you see a red light, you think "stop" — we have what I will call an external thought. Its origin is outside the mind, even though the association happens inside it.

```
sensorial stimuli  →  direct association  =  external thought
```

The active mind

Behind every thought sits an active mind. By active mind I mean the conscious working space in which thinking happens: short-term memory, plus the parts of long-term memory currently in play, minus the unconscious processes that the autonomous nervous system runs in the background to keep us breathing, balanced, and alive. The active mind is what you can introspect on. The unconscious is what you cannot.

```
active mind  =  short-term memory + long-term memory – unconscious
```

Internal thought

An internal thought — what we ordinarily call "my own idea" — is what happens when the active mind chains associations across stored memory without a fresh sensorial trigger. You are walking, no one has spoken to you, nothing has changed in your environment, and suddenly you remember you owe a friend a phone call. Where did that come from? From you. From the active mind associating one item in memory with another.

```
active mind  +  (association × n)  =  internal thought
```

Useful thought

Most internal thought is brief, fragmentary, and not particularly useful. It is the mental equivalent of changing television channels. A useful thought — what we

sometimes call rational thought — is something more demanding. It is a sequence: a chain of associations that link existing structured memory items in a new way and produce something the mind did not already contain. A solution. An insight. A plan.

```
useful thought  =  ((structured memory + association) ×
                     n_sequential)
```

The clearest way to picture this is the metaphor of a train. Each carriage is a memory item already in storage — a fact, a concept, a recollection, a procedure. Useful thought is the new track that links the carriages together into an order that did not exist before. The carriages are old; the track is new. That is the work.

Two consequences follow from this model, and both will matter later.

First, useful thought is sequential. It does not happen all at once. There is always an order in which the associations form, and that order is itself part of the structure of the thought. You cannot skip steps; you can only appear to, by jumping to a conclusion using a shortcut your mind has accepted as a substitute for the missing steps.

Second, useful thought is slow. Logic is a branch of mathematics, and mathematics is exact, and exactness has a price: time. The active mind compensates by taking unconscious shortcuts. Most of the time, these shortcuts are good enough. Sometimes they are not. Sometimes they are catastrophically wrong. Cognitive psychology has documented roughly two hundred such systematic shortcuts; the canonical inventories are widely available, and most of the manipulation techniques described later in this book correspond, term for term, to one or another of them. Each one of them, in principle, can be weaponized into a manipulation method.

Cognitive psychology calls these shortcuts heuristics or biases. From the perspective of the person being targeted, they are something simpler and more dangerous: doors that someone else might know how to open.

1.2 A Note on Rational Thought

Rational thought is governed by absolute logic. A statement — any statement, including any thought — is true, false, or unknown. There are no other

categories. Rational thought assumes full trust in one's own mind. It makes an abstraction from any factor that might cause doubt.

This does not mean the doubt is gone. It means the doubt is not factored in during reasoning. You make a footnote about it in structured memory — "this premise is shaky" — and you proceed. If you tried to reason while continuously second-guessing every premise, you would never reach a conclusion at all. Doubt is supposed to be a marker, not a brake.

This is why a manipulation campaign that succeeds in destabilizing your trust in your own mind is so devastating. It does not have to make you believe anything specific. It only has to make you unable to make abstractions — to remove the brake from off and force it permanently on. Reasoning grinds to a halt. You can still feel; you can still react; you cannot any longer think things through to the end. The mind that cannot complete a thought is, for all practical purposes, the mind of a person who has been disabled.

1.3 Where the Attack Happens

With the model in place, we can describe the basic shape of psychological manipulation as it is experienced by a Targeted Individual. The attacker's strategic objective is almost never to insert a single false belief. The strategic objective is to misplace trust — to move trust away from the target's own mind and onto the attacker, the situation, or the noise. Once trust has migrated, specific beliefs can be inserted at will.

There are several mechanisms by which this misplacement is achieved. They are not mutually exclusive; in any sustained campaign they tend to be layered.

Confusing the target's own thoughts with external suggestion

The first method is the most elegant, and in some ways the most cruel. It exploits the fact that the boundary between conversation and thought, in memory, is thinner than people assume.

Consider an everyday example. You are working on a piece of code. Your mind drifts; you find yourself daydreaming about something else. Suddenly you notice an error in your code — the kind of structural bug that would have caused an infinite loop, the kind you might never have noticed if you had stayed focused

on the surface task. The realization arrives clean, with the satisfying click of a piece falling into place.

Where did that realization come from? From you. From your active mind associating, in the background, a piece of structured knowledge about your own scripts with a half-remembered concern about file extensions. It was a useful thought, exactly as defined above.

Now suppose someone you do not entirely trust steps in and claims, casually, that this insight was given to you. "That error," they say, "I told you about that. Don't you remember?" They didn't. But the conversation is plausible enough — it could have happened, you have had hundreds of similar conversations — and you cannot positively prove that it did not. Your memory of having had a conversation with this person is foggy, as memories of casual conversations always are. Their certainty is high. Yours is not.

If you accept their account even tentatively, something has shifted. The attacker has become, by your own concession, a source of useful thought. And once a person becomes a source of useful thought, the mind starts to defer to them on adjacent matters. The trust has begun to migrate.

Counter it by remembering this: there is no real answer to the question "where do thoughts come from?" If an idea emerges in you without a direct sensorial trigger — if no one has just spoken to you, no one has just shown you something — then by definition it is your own idea or thought. Even if previous input from senses has contributed to the structured memory items the thought is built from, the act of association that produced the new thought happened in you. There is no other place it could have happened.

Useful thoughts that do not come from sensorial stimuli, and that are rational, have no other possible source than your own mind. This is not a comforting slogan; it is a tautology. Hold on to it.

Exploiting structure to fabricate memory

Conversations leave specific traces in memory. They have structure: who said what, when, where, in what tone. When a manipulator wants to claim a conversation that did not happen, the giveaway is usually that the structure is missing. They will give you the content but not the context — the "what" without the "where" and "when."

If a person tells you, “We talked about this,” and you cannot place the conversation — not the room, not the day, not the lighting, not the words immediately before or after — the most likely explanation is that the conversation did not occur. The mind, given a content claim without structural anchors, will sometimes manufacture the missing structure to keep the social interaction smooth. Resist this. The absence of structure is itself information.

Speeding up the thinking to force shortcuts

A second method targets the slow nature of rational thought directly. If rational reasoning is slow, then anything that forces faster decisions pushes the target out of rational reasoning and into shortcuts. Stress does this. So does urgency. So does emotional flooding. So does sleep deprivation. So does the steady, low-grade exhaustion of never being able to relax — because someone or something is always, somehow, intruding.

Under speed, the mind defaults to its catalogue of unconscious shortcuts — those thousand-plus shortcuts catalogued by cognitive science and weaponized by the more cynical end of the marketing and political-influence industries. A target whose mind has been forced into permanent shortcut mode is a target who can be steered.

The countermeasure is structural, not heroic. It is not to think harder — you cannot outwork an attacker who has time on their side. It is to reduce, wherever possible, the conditions that force speed. Sleep. Eat. Make notes. Slow down. Refuse to make any important decision under acute stress. The single most effective defense against shortcut-based manipulation is not intelligence; it is rest.

1.4 The Pattern of a Manipulative Conversation

Manipulative conversations have a recognizable arc. Once you have seen it a few times, you can see it coming. I describe it here in its basic form because the rest of this book will return to it repeatedly.

The arc has four stages. They do not always appear in this order, and not every stage appears in every conversation, but the underlying logic is consistent.

Stage one: open with curiosity and warmth

The conversation begins with the manipulator showing apparent interest. They ask about something specific: a project, a name, a view of yours. The interest seems genuine, and may in fact be partly genuine — manipulators are often curious about their targets in the way that an entomologist is curious about an interesting beetle. The target, who has no reason to expect hostility and is naturally inclined to give a person the benefit of the doubt, opens up.

This stage exploits something I think of as the survival default of social interaction. We tend to give unknown people a certain quantum of trust. It is not foolishness; it is what makes society function. A target with a healthy disposition will extend that trust automatically, especially if they retain, by choice or by nature, some degree of naïveté about other people. (Many targets are naive on purpose. Cynicism, fully embraced, is its own kind of death.)

Stage two: reinforce goodwill

Once the target is engaged, the manipulator expresses sympathy. They share the target's view. They offer a small insight — something that seems useful, relevant, perhaps even constructive. This phase produces the sensation that we have all felt and trusted: "Finally, someone who understands."

The function of this stage is not friendship. The function is to establish the manipulator as a benign, even beneficial, presence — so that whatever follows can be framed as coming from a friend. The same words said by an enemy and a friend land very differently in memory. The manipulator wants to be filed under "friend" before they say the things they actually came to say.

Stage three: confuse and belittle

The third stage is where the conversation's real purpose appears. A second person may join in — a distractor whose role is to interrupt, redirect, talk over, or introduce non-sequiturs. The conversation begins to lose its thread. You find yourself losing track of what point you were making. You explain yourself; you are misunderstood; you explain yourself again; you are misunderstood differently. Eventually, after enough confusion, comes the belittlement. A small, almost-cherishable phrase from the manipulator — the tone of an indulgent adult correcting a slow child.

By the end of this stage, the target has been positioned as the source of the trouble. The blame for the conversation's failure lands on them. They were unclear. They were confused. They were too much.

The technical vocabulary for this is gaslighting, but I find that word has been so widely used in popular discourse that it has lost most of its descriptive precision. What I want to name is something more specific: the deliberate, layered combination of confusion-induction and status-lowering that, when repeated, hollows out a target's ability to trust their own perception of any conversation.

Stage four: raise hope, then crush it

The cruelest stage, and the most diagnostic, is the last. After the belittlement, the manipulator may relent. The conversation, against expectation, turns warm again. The target — starved of normal human interaction by months or years of this treatment — feels hope rise. Perhaps this person will, after all, be reasonable. Perhaps an understanding can be reached.

And then, with timing too precise to be accidental, the hope is beaten down. A small cruelty, a casual dismissal, a sudden withdrawal. The target is left wondering: what was that for?

There is a specific reason for this final move. A target whose hope has only ever been disappointed eventually stops hoping; that target is hard to reach. A target whose hope is repeatedly raised and then crushed remains permanently hopeful and permanently devastated, an emotional state that is far more disabling than simple despair. The instability is the point.

If you recognize this arc — the warmth, the sympathy, the confusion, the belittlement, the hope raised and crushed — you are looking at a manipulation pattern, not a difficult relationship. Naming it does not end it. But it stops you blaming yourself for it.

1.5 Brainwashing, Reconsidered

The word brainwashing carries unhelpful associations. It evokes Cold War clichés, cults, and mind-control science fiction — categories that mainstream psychology has, with some justification, treated skeptically. The term has become an obstacle to the very phenomenon it was meant to describe.

I want to redefine it here, narrowly and operationally.

Brainwashing, as I will use the term, is not primarily an unconscious phenomenon. It is something the target remembers but cannot correctly

evaluate. The events happened; the target can recall that they happened; what the target cannot do is assemble those memories into a coherent assessment of what the events meant. The mind keeps circling the same period of time, sensing wrongness, never quite resolving it. The victim, in effect, knows something is wrong but cannot finish the sentence that says what.

Why? Because the victim was deprived, during the period of conditioning, of the conditions that allow correct evaluation. They were stressed. They were isolated. They were sleep-deprived. They were dissociated, in the technical sense — cut off from their normal sense of continuous selfhood. Above all, they were given pieces of a reality and not the picture those pieces fit into. The pieces themselves were factual, or near enough; it was the assembly instructions that were missing or false.

If you assemble the pieces wrong, the picture you produce is the one the manipulator wanted you to produce. And because the pieces are real — because you genuinely remember each individual piece — the false picture feels grounded in fact. It is not a hallucination. It is something worse: a coherent story that fits the available evidence, and that has been engineered to fit.

From this redefinition, several practical consequences follow.

- Brainwashing is reversible, in principle. What was done by manipulating the assembly of the pieces can be undone by re-examining the assembly. This is hard, but it is not magic.
- The victim must learn what went wrong, not be told. Being told the conclusion does not help; it adds another piece to the same broken puzzle. The victim has to walk back through the pieces themselves and see how they fit differently.
- Reprogramming is incomplete without a new purpose. A mind that has had its old purpose hollowed out cannot stay empty; it will fill with something. If the victim does not provide a new purpose for themselves, the manipulator's purpose tends to flow back in by default. Recovery requires building, not just deconstructing.

1.6 The One Thing That Cannot Be Manipulated

I will end this chapter on a note of structural optimism, because the model has one last feature that the rest of the book depends on.

The active mind can be manipulated. Memory can be confused. Trust can be misdirected. Hope can be raised and crushed. Almost everything in the architecture I have described is, under sustained pressure, vulnerable.

Almost everything. There is one exception.

Beneath the active mind, beneath rational thought, beneath the catalogue of cognitive shortcuts, sits something older: the survival instinct. It is the fastest part of us — faster than language, faster than reasoning. When a real threat appears — not a thought of a threat, not a story about a threat, but the actual thing — the survival instinct fires before the conscious mind has finished forming the question. Adrenaline arrives. Attention narrows. The mind analyzes first and then acts.

I define the survival instinct, for our purposes, very simply:

```
survival instinct  =  curiosity  +  distrust
```

Curiosity, because the survival instinct must understand what it is dealing with before it can act. Distrust, because under threat the default has to be: not safe until proven safe. Together, those two impulses have kept our species alive for the entire span of its history. They are older than civilization, older than language, almost certainly older than self-awareness as we understand it. They cannot be argued out of you. They cannot be socialized out of you. They cannot, so far as I have been able to determine, be manipulated.

The reason they cannot be manipulated is that they do not run on the same machinery as conscious thought. They run, instead, on something deeper — something that does not consult the mind for permission. The manipulator, however skilled, is operating on the layer above. The survival instinct sits below, and watches.

If you have read this far and recognized any part of the patterns I have described, then I will tell you something I believe is true and that I have staked a significant portion of my own life on: the part of you that recognized them is the survival instinct doing its job. It is the part of you that has been telling you for a long time that something is wrong, even when every part of the surrounding world insists that everything is fine. It has been right. It will continue to be right.

The work of the rest of this book is to give that instinct the vocabulary and the framework it has been waiting for — so that the alarm it has been sounding can finally be translated into something the rational mind can act on.

We begin there.

CHAPTER 2

The Architecture of Doubt

How Belief, Knowledge, and Uncertainty Get Weaponized — and How to Build Them Back

---

"The moment faith starts to contradict reality, it will inevitably lead to disappointment and avoidable existential questions."

— Author's notebook, January 2023

We have spent two chapters describing what manipulation does to a target. We have described the architecture of thought, the arc of a manipulative conversation, the moment the audience joins, the descent from logic into bullying, the trauma in the place where abuse was forgotten, the geometry of resilience, and the wall of the will to change. What we have not yet examined directly is the single most important psychological resource that a manipulation campaign tries to take from you, and the one it returns last during recovery.

That resource is your relationship with what you know.

Not what you believe. Not what you hope. What you actually know — the small, hard floor of facts about your own life, your own perceptions, your own experience, that nothing said to you can shake. Manipulation campaigns target this floor specifically. They do not, mostly, try to convince you of false things; that would be too obvious and too easily refuted. They try, instead, to make you uncertain about true things. The asymmetry is everything. A confident lie is fragile; a manufactured doubt is durable. This chapter is about why doubt is the manipulator's most valuable instrument, what it costs you to be the one carrying it, and how the relationship between believing, knowing, and doubting can be rebuilt on terms that you set.

2.1 Why Doubt Is the Manipulator's Most Valuable Tool

There is a clean way to see why doubt, specifically, is the asset a sustained manipulation campaign cares most about preserving. It is a thought experiment in two parts.

Imagine, for a moment, that you had less doubt about what was being done to you. Suppose you knew, with the kind of stable certainty that ordinary people have about ordinary facts, who was responsible, what their motives were, what your own role was. Even without a single new piece of physical evidence, your situation would have changed substantially. You would gain leverage. You could plan. You could choose where to push and where to wait. You could explain yourself to others without the constant qualifying of your own claims that doubt forces on you. You could sleep differently. The volume of your decisions would expand.

Now imagine the opposite: that the doubt has been carefully maintained, even cultivated, over years. You can no longer trust your own observations; every fact you notice is shadowed by the question of whether you are noticing it correctly. You cannot plan, because you are not sure what you are planning against. You cannot explain yourself to others, because before any explanation reaches them, you have already qualified it three ways. You cannot rest, because resting requires a settled mind, and the doubt has unsettled the very settling.

The first state is dangerous to the manipulator. The second is the manipulator's ideal.

This is why — and the realization, when it lands, reorders many things — the campaign's effort is not, primarily, directed at making you believe specific false claims. It is directed at preventing you from arriving at stable beliefs at all. The methods are familiar from the previous chapters: the carefully timed contradiction, the question-instead-of-answer, the introduction of plausible alternative explanations whenever you reach a conclusion, the quiet undermining of any source you might have used to verify. None of these techniques requires you to accept a particular falsehood. They only require that you not arrive at a particular truth.

I will call this strategy of keeping a target in suspended judgment a disruption strategy, with the caveat that the term is descriptive rather than drawn from any settled clinical or technical literature. The strategic logic is that a target unable to consolidate beliefs is a target unable to act. Action requires the courage of a position. A person held permanently in suspended judgment

cannot get to courage, because they cannot get to a position. Their distress is real, their analysis may even be accurate, but the distress and the analysis cannot become motion.

Maintaining doubt in the victim, then, is what gives the victim very few means to resist or fight back. It is the lock on the door. Recognizing that the lock exists is the first step toward picking it.

2.2 What You Are Not Allowed to Be Sure Of

If you watch the campaign carefully — and most targeted individuals eventually do, with the painful precision that comes from having no other defense — you can identify the specific things you are not being allowed to know. They are usually two.

First, you are not allowed to know who is doing this. The cast of characters remains permanently fluid. Every time you converge on an identification, a new piece of misdirection is introduced. The known suspects are revealed to have been unaware. New suspects are floated and then withdrawn. Sometimes the question is reframed entirely — perhaps it is not a person at all, perhaps it is a system, perhaps the question itself is incorrect. The shape of the perpetrator is kept just out of focus, indefinitely.

Second, you are not allowed to know what your own role is. Are you a victim, or are you in some sense participating? Could it be that some of what is happening is your own doing, your own failure, your own error? Is the conviction that you are a target itself a symptom of something? Each time you arrive at stable ground here, the ground is moved. The campaign wants you uncertain about your own position in the situation, because a person uncertain about their own position cannot make a reliable case to anyone, including themselves.

These two uncertainties — about who and about you — are not, primarily, intellectual. They are operational. As long as both remain unsettled, you cannot effectively report what is happening to anyone whose help you would need: a doctor, a lawyer, a friend, a journalist, an investigator. Each of these professionals requires from you a stable, declarable account: this is what is happening, this is who is doing it, this is my role. Without that account, they cannot help you, however much they might want to. The campaign is, in this sense, less about what is being done to you and more about preventing the kind of account that would let someone else come to your aid.

When you see this clearly — when you see that the disruption of your account is the operation, not a side effect of it — you can begin to refuse to participate. The refusal does not have to be loud or dramatic. It is mostly internal. You stop revising your account every time a new doubt is introduced. You write down what you actually know, distinct from what you have been made to wonder about. You hold the floor. You do not let the floor be moved by every gust of new ambiguity.

Holding the floor is harder than it sounds. The doubts that have been planted have, over time, become indistinguishable from your own honest skepticism, which is one of your better qualities. You will need a way to tell them apart. The next sections give you the tools.

2.3 Don't Entertain the Excuses

Here is a small rule that will save you years if you learn it now.

When someone who has hurt you offers you an explanation for why they hurt you, your instinct will be to receive it as data. To weigh it. To see whether it is plausible. To incorporate it into your understanding. This is what reasonable people do with information. You are a reasonable person. You will do it. Almost everyone does.

Stop doing it.

Excuses for abuse are not data; they are part of the abuse. Their function is not to inform you. Their function is to occupy your mental processing time with ambiguity that benefits the person offering the excuse. While you are weighing whether the excuse is true, you are not pursuing the consequences of the underlying behavior. The excuse buys time. Sometimes the time it buys is everything.

The pattern is often quite specific. A series of manipulative excuses are offered, each one with enough surface plausibility that a generous mind can entertain it. The cumulative effect is what one of my own notebooks calls the romantic-comedy intercession — a tonal shift that frames the cruelty as endearing, or as misunderstanding, or as a quirk to be smiled at. A note to myself from January 2023, written after one of these episodes, said simply: do not entertain their excuses for their abuse. I underlined it. I am underlining it again here, because three years later it still requires the underlining.

This is not a rule against all explanation. People who hurt others sometimes do come to genuine understanding of what they did and why; that understanding can be valuable, in the right context, after enough time, with enough accountability. But the offering of excuses in the middle of an ongoing pattern of harm is a different thing. It is a tactical move. It does not deserve the courtesy of being treated as if it were sincere.

The simple rule: while the abuse is still in progress, the explanations for it are also abuse. Treat them accordingly. You are allowed to refuse to discuss them. You are allowed to say, if asked, that you will consider the explanations after the behavior has stopped for a verifiable amount of time. You are allowed, most importantly, not to be the person on whom the burden of evaluating these explanations is placed. That burden does not belong to the target. It has been transferred there illegitimately, and you can transfer it back simply by declining to carry it.

2.4 The Perfectionism Trap

There is a particular variety of internalized doubt that targeted individuals tend to develop, and that I want to name carefully because it operates from the inside and is therefore harder to see than the externally imposed kinds.

It is perfectionism, in a specific weaponized form.

The pattern goes like this. Inside you — not as a voice from outside, not as an external accusation, but as the texture of your own thinking — there is an emphasis on what you have not yet done. On the corner that is not yet finished. On the email that has not yet been answered. On the project that is not yet shipped. On the relationship that is not yet repaired. The completed work, however substantial, recedes. The incomplete work, however small, becomes the foreground of your attention. Over time, this becomes the default shape of your self-evaluation: you are always behind. You are always failing. The list of things you have not done is the only list that registers.

The clinical literature would call this a depressive cognitive distortion. That naming is correct but, I think, incomplete. In the case of a targeted individual, the distortion has often been deliberately reinforced by external sources — by reminders of what you have not done, by emphases on any weakness, by the recasting of all your incomplete work into failure. After enough rounds of this,

the external voice and the internal voice converge. You no longer need anyone to tell you that you are failing; you tell yourself, in their voice, all day, every day.

The cost of this is severe and non-obvious. Perfectionism, in this form, almost guarantees that you only see failure. Every successful piece of work is invisible, because the test by which it is being evaluated is not whether it succeeded but whether it is finished and unimprovable, and almost nothing meets that test. You build something genuinely valuable; the perfectionist eye sees only the rough edges. You complete a difficult task; the perfectionist eye sees only the next difficult task that has now appeared. The volume of your achievement does not rise in your felt sense of yourself, no matter how much you objectively achieve.

Here is the strange thing about this trap. It looks, from the outside, like high standards. It looks, from inside, like high standards. But it does not function like high standards. Real high standards are calibrated to specific work, applied to specific work, and then released when the work is done. Weaponized perfectionism is calibrated to nothing in particular and applied to everything indiscriminately. It does not raise the quality of your work; it lowers your capacity to recognize that the work was ever good. It is, in its effect, a tool for self-discouragement that wears the costume of self-improvement.

Three practical moves help with this, and I offer them as habits to be built, not as a one-session fix.

- First, keep an explicit list of completed work. Not a to-do list — the to-do list is the perfectionist eye's preferred document, and it never empties. A done list. What did you actually finish today, this week, this month? Read it before you read anything else. The completed-work list is the corrective lens.
- Second, set the criterion for done in advance, before starting the work, not retroactively. Perfectionism is the practice of moving the goalposts after the play. If you set the goalposts before the play and you cleared them, the play succeeded. The fact that other, taller goalposts now exist further down the field does not retroactively undo the success of the original play.
- Third, notice when the perfectionist voice is doing the speaking. It has a recognizable register. It is concerned with what is missing. It frames partial work as failed work. It compares your output to ideals rather than

to alternatives. When you notice it, you do not have to argue with it. You only have to refuse to be its student. Listen, register, and continue.

2.5 Believing, Knowing, and the Difference Between Them

I want now to go to the philosophical center of this chapter, because the practical work of the previous sections rests on a distinction that ordinary language tends to blur.

Belief is what we hold about the world when knowledge is not available. We believe a friend is trustworthy because we have not yet had the experience that would settle the matter; we believe the train will arrive on time because we have not yet looked at the platform. Belief is a placeholder. It fills in for knowledge in regions where knowledge has not yet reached, and it functions, when it is healthy, as a kind of rough draft of what we expect to find when we look more closely.

Knowledge is different. Knowledge is what we have when we have looked. It is more constrained — we know fewer things than we believe — but the things we know have a different texture. They do not require permanent maintenance. They are stable across moods. They do not need to be defended against contradiction, because they are simply what is.

Doubt sits between the two. Doubt is what tells us that a belief is not yet knowledge. It is, in this sense, a tool of the mind — the part of cognition that says: not yet, look further, you do not have enough to call this settled.

In a healthy mind these three sit in a workable triangle. Belief and knowledge are the two stable poles; doubt is the working motion between them. New experience and new evidence move beliefs into knowledge, or sometimes — humbly — demote former knowledge back to belief when something we thought we had seen turns out to have been wrong. Doubt is the engine that drives this circulation. It is the friction that produces motion.

Belief, by itself, has a particular use that is worth naming. Belief is a way to avoid desperation. When the future is uncertain in a frightening way, belief allows you to live forward into it without collapsing. The terminally ill patient who still hopes for a miracle is using belief in this stabilizing function, and we should not begrudge them this; the alternative is despair, and despair is a kind

of dying that arrives before death itself. Belief, in this register, is an anti-desperation technology, and human beings need it more than they like to admit.

But belief has a cost. It is a way of running away from cold facts. If the facts are temporarily intolerable, this is fine — a temporary withdrawal can be the difference between collapse and continuation. If the withdrawal becomes permanent, however, belief slides into avoidance, and avoidance prevents the action that might still be possible. The terminally ill patient who only believes in miracles, and never investigates the actual treatments available, has used belief well in the first hour and badly in the long run. The same belief, depending on what it allows or prevents, is medicine or poison.

There is a passage in my notebooks that I want to quote almost verbatim because I cannot improve on what I wrote at the time. Believing, I wrote, is a way to not have to make the effort to try to understand and know for further (and better-informed) experiences that allow personal growth. That is harsh. It is also, often, true. Belief used as a substitute for knowledge prevents the growth that knowledge would have made possible.

The healthy alternative is not to abolish belief; that would be cruel and impossible. The healthy alternative is what I will call critical understanding: the practice of letting belief, faith, and conviction continue to operate, while also continuing to test them against current reality and against what is genuinely known. Faith and believing can coexist with doubt, on this picture, as long as the doubt is allowed to do its work. The moment faith starts to contradict reality, however, it will inevitably lead to disappointment and to the avoidable existential questions that come when reality refuses to behave the way faith requires.

What you want, in other words, is faith that is responsive — that updates when reality presents new information, that does not need to be defended against the world, that does not require you to make yourself ignorant in order to keep it intact. The practice of holding faith and doubt simultaneously is not a contradiction. It is the actual structure of a working mind.

2.6 When the Insurmountable Has to Be Faced

There is a particular situation that targeted individuals encounter often, and that the framework above lets us discuss with more precision than we usually manage.

The situation is this: you are facing something that appears insurmountable. The problem is too large; the resources are too few; the time is too short; the bridge across the chasm cannot, on the available evidence, actually be built. The temptation, in this situation, is to retreat into belief — to convince yourself that some solution will appear, that some help will arrive, that the situation is not, after all, as bad as the cold facts suggest. This retreat is not contemptible. It is a survival reflex. But it is not, by itself, sufficient.

The cold-facts-only response is also not sufficient. If you stand at the edge of the chasm and only count the meters, refusing to act because no bridge yet exists, you will surely perish. The pure rationalist position, taken alone, is as much a failure mode as the pure believer position. Both lead to the same place, by different routes.

What works, instead, is an alternation. First, look closely at the insurmountable. Analyze it. See its actual shape. The temptation to look away — because looking is painful — has to be resisted. You may, in the close looking, see a way out that was not visible at first; many "insurmountable" problems contain a small unguarded passage that becomes visible only when you stop flinching from the larger structure. If you do not see a way out, the close looking still serves a purpose: it lets you accept the reality of what you are facing, including, in extreme cases, your own ending.

Then, having looked, you can use belief deliberately. Not as a substitute for the analysis, but as a continuation of it. You can choose, knowing what you know, to act as if a bridge can still be built. The choice is now informed; it is no longer escapism. It is the deployment of belief as a working assumption, on top of analysis that has not been skipped. Many of the human achievements we admire most were built this way: someone looked clearly at the cliff, found no bridge, and started building one anyway.

Belief, in this disciplined form, becomes a way to extend reach beyond what is currently provable. It is no longer the avoidance of reality. It is the use of reality as a launching surface. The same word — belief — names two very different practices, and only one of them is helpful. The difference between them is whether the analysis came first.

2.7 The Three Components You Cannot Live Without

Before closing this chapter I want to offer one more model, the third in this book, because the architecture of recovery becomes much easier to plan once the architecture of the self is sketched a little more fully.

Three components, I believe, make up the core of an ordinary functioning person. They are not the only things a person consists of, but they are the three that cannot be fully absent for long without serious damage to the others.

- Sensitivity is the capacity to be affected by the world — by other people's states, by beauty, by suffering, by texture, by tone. It is the channel through which experience reaches you. A person without sensitivity is not callous, exactly; they are simply not in contact. Things happen and the things do not register.
- Intelligence — here in the broad sense, including reason but not limited to it — is the capacity to make sense of what sensitivity has delivered. To pattern-match across experiences, to test hypotheses, to predict, to plan. Intelligence is what turns the raw incoming material of sensitivity into useful structure.
- Empathy is the capacity to recognize that other minds, structurally similar to your own, are also receiving and processing experience. It is the bridge between your inner life and the inner lives of others. Without it, the other people in your world are not quite real to you; they are figures in a landscape that contains only your own felt life.

The interesting structural fact about these three is that they cannot exist for long without each other. Sensitivity without intelligence becomes overwhelm — the person is affected by everything and able to make sense of nothing, and the result is paralysis or breakdown. Intelligence without sensitivity becomes a kind of clever sociopathy — a mind that processes patterns without feeling the substance of what those patterns describe. Either of these without empathy yields cruelty, of the active or the passive variety: the world is full of objects to be analyzed or felt, but no one is in the world with you.

What targeted abuse damages, characteristically, is not one of these in isolation but the connections among them. The sensitivity is still there but no longer trusted, because too much of what it picked up turned out to have been engineered. The intelligence is still there but no longer applied, because every conclusion it reaches gets undercut. The empathy is still there but increasingly costly to feel, because feeling for others requires a stable base of feeling for oneself, and that base has been eroded.

Recovery, in this picture, is partly the slow re-coupling of these three. You learn to trust your sensitivity again, by paying attention to small private impressions that you do not have to defend to anyone. You let your intelligence work again, by allowing yourself to reach conclusions and hold them, rather than disqualifying every conclusion as soon as it arrives. You let your empathy come back online, slowly, with people who have shown over time that they are safe places for it. None of these moves is dramatic. All of them are necessary. The figure rebuilds in the connections, not in the components.

2.8 What You Already Know

I want to end this chapter with the observation that has, in the long run, mattered more to me than any other I have written down in five years of notebooks.

If you have been targeted for any sustained period, there are things you already know. You know them despite the doubt that has been planted around them. You know them despite the manipulators' best efforts to make you not-quite-know them. You know them in the way that people know the smell of their childhood home, or the sound of a parent's voice — not as conclusions arrived at by reasoning, but as facts that are simply there, underneath.

You know that something has been happening to you that should not have been happening. You know roughly when it started, even if the start date is fuzzy. You know that the explanations you have been offered for it do not, when you sit with them honestly, hold together. You know which people have been kind to you in ways that proved real, and which people have been kind in ways that did not. You know what your own conduct has been, in the central things; you know where you have made mistakes and where you have not.

These knowings are a small inventory. They will feel inadequate — surely there should be more, given how much has happened? — but the inventory is what you have, and it is more than you think. The disruption strategy of the last chapter has been working hard to keep you from looking directly at this inventory. It wants you focused on the missing pieces, on the unresolved questions, on the ambiguities. It does not want you to count what is actually settled.

Count it anyway. Write it down, if writing it down helps. The inventory of what you actually know, separated from what you have been made to wonder about,

is the floor you have been told for years that you do not have. You have it. Most of the rest of recovery is the work of standing on it.

We will turn, in Chapter 4, to the practical question of what a person with a stable inventory of their own knowledge can begin to do — not in the abstract, but in the small, recoverable, daily ways that the work of getting one's life back actually proceeds. The big gestures, in my experience, are rarely the decisive ones. The decisive moves are smaller and more numerous than the recovery literature usually admits. We will look at them carefully.

But for now, this chapter ends where it began: with the observation that doubt is the lock, and that the lock is more pickable than it looks once you have seen it for what it is. The keys are unspectacular. They are patient. They work.

❦ ❦ ❦

PART II

The Methods of Manipulation

How a Person's Thinking Is Attacked, in Detail

CHAPTER 3

The Shock That Does Not Land

On Trauma, Sleep, and the Mechanics of Healing

"Trauma grows. If the shock response is not used,
the healing of the trauma disappears."

— Author's notebook, November 2022

I have promised, three times now, a chapter on the practical and unspectacular daily moves that the work of recovery actually consists of. I am, in this chapter, finally going to deliver part of that promise. But not in the way the recovery literature would lead you to expect. The standard advice — eat well, sleep enough, exercise, stay connected — is correct as far as it goes, but it does not go far enough to be useful to a person whose situation is the one this book has been describing. The standard advice assumes that the body and the mind are operating in conditions where ordinary maintenance is sufficient. The conditions a targeted individual lives in are not ordinary, and the maintenance has to be calibrated to those non-ordinary conditions if it is going to do any good.

This chapter is built around a single observation that, once seen, organizes the rest of the practical material: trauma, in a sustained manipulation campaign, does not heal on its own. It does not heal slowly. It does not even heal badly. It does not heal at all, because the campaign is structured precisely to prevent the conditions under which healing happens. Recovery, in such conditions, is not a matter of waiting for time to pass. It is a matter of restoring, by deliberate effort, the conditions that the campaign has removed. Once you understand which conditions those are, the practical moves follow.

3.1 Why the Trauma Does Not Heal

Begin with what ordinary trauma processing looks like, when nothing is interfering. A person experiences a difficult event. The event registers on the body and the mind in several layers at once: a sensation, an emotion, a thought, a behavioral impulse. There is, immediately afterward, a period in which the

system is processing — often poorly, often confusingly, often with the appearance of dysfunction. This period can look from the outside like the person is not coping. It is not that. It is the system doing the very specific work of integrating an event that has not yet been integrated.

During this period, the person typically goes through a sequence that, when allowed to run to its end, produces healing. They register the event again, in memory. They have a shock response — emotional, somatic, sometimes embarrassingly intense, sometimes delayed. They reach a point of recognition: this happened, it was bad, I am still here. They reorient. The event becomes a thing that occurred, with a beginning and an end, that they survived. The trauma does not vanish; the relationship to it becomes processable. The person can speak about it, even if not always comfortably. They can recall it without re-entering it.

There is a sequence here that is widely recognized in the trauma literature, even if it is not always given a single name: a catharsis that the body needs to undergo in the wake of the precipitating event. The catharsis is the shock response and what follows it. Without that catharsis, the integration does not happen. The event remains unprocessed. The body continues to carry the unfinished response.

In a sustained manipulation campaign, the catharsis is, very specifically, what is being prevented. The campaign does not prevent events from occurring; events occur, often more than the target can keep up with. The campaign prevents the events from being followed by the response that would integrate them. There is no quiet period afterward in which the body can catch up. There is no space in which the shock can land and complete its work. The next event arrives before the previous one has been processed, and the processing is interrupted, and the unprocessed material accumulates.

This is why the trauma of a targeted individual feels, even years in, somehow fresh — somehow as if it is still happening, even when the specific event one is thinking of took place long ago. It is still happening, in the operational sense. The integration has not occurred. The shock response that would have closed it has been prevented from completing. From the body's point of view, the event has not yet ended, because the work that would have ended it has not been done.

This reframing has a useful consequence. It locates the work of recovery, in this domain, in something quite specific: restoring, deliberately, the conditions under which delayed shock responses can finally land. The work is not to think one's way out of the trauma. It is to allow the body to do, belatedly, what it would have done years ago if it had been allowed to.

3.2 The Three-Step Loop

Let me give the structure of trauma processing in its clearest form, as a three-step loop. Each step has to occur for the loop to complete. If any one of them is interrupted, the trauma stays in the body unprocessed. The three steps are these.

First, *memory of the event*. The person re-encounters the event in memory. This re-encounter has to be specific enough to anchor what the body is going to respond to. A vague sense that something bad happened is not sufficient. The mind has to bring forward the event with enough detail that the body knows what is being remembered. This is why generalized recovery techniques that try to soothe trauma without naming it often produce only partial results — the body cannot release what it has not been told to address.

Second, *shock response*. The body, having been told what is being remembered, produces the response it would have produced at the time, if it had been allowed. This response is often physically intense. There may be tears that arrive without an immediate trigger; there may be tremor; there may be heat or cold; there may be a kind of grief that does not seem to belong to the present moment because it does not. It belongs to the moment of the original event. It is arriving late, but it is arriving correctly. The shock response is the body finally completing the work that was interrupted.

Third, *reality and correction*. After the shock response has run its course, there is a period of reorientation. The person notices, often with surprise, that the event is over. They are still here. The world is still recognizable. The relationship between the past event and the present moment is re-established correctly: the event was then, this is now, those are different. The integration has occurred. The trauma has not been erased; it has been moved into a place from which the rest of life can resume around it.

All three steps have to happen for the loop to close. If only the first happens — memory without shock response — the trauma is repeatedly recalled without

ever being processed, which is its own kind of damage. If the first and second happen but not the third — shock response without subsequent correction — the person is left in the response itself, sometimes for hours or days, with no orientation back to the present. If the second is repeatedly interrupted before it completes, the body learns that shock responses do not end, and stops initiating them at all. This last failure mode is, I believe, the one most targeted individuals end up in. The body has stopped trying to complete the loop, because every previous attempt was interrupted before completion. The trauma sits there, durable and unintegrated, behind a body that no longer knows how to address it directly.

Recovery, on this picture, is the slow re-permission of the loop. The body has to learn, one event at a time, that the shock response is now allowed to complete. This is why recovery works best in conditions of safety, with someone present whose presence is reliable, in environments that do not interrupt. The conditions are not optional. They are the conditions under which the body can risk a shock response without expecting it to be cut off again.

3.3 Overload as a Method

Once the loop is named, the methods by which a campaign prevents it from completing become easy to describe. The most reliable of these methods is the one I will call *overload*.

A target experiences an event. The event is sufficient, by itself, to warrant a shock response. The body begins to organize one. Before the response can complete, a second event arrives. It may be unrelated to the first; what matters is that it is sufficiently demanding of attention that the body has to redirect its resources toward dealing with it. The shock response to the first event is interrupted. Then a third event arrives, and a fourth. The target's available bandwidth for response is, at any given moment, fully consumed by whatever is currently happening. Nothing ever gets to settle into the second step of the loop.

From the outside, this looks like the target is having an unusually difficult life. From the inside, it is something more specific than that: it is a structural prevention of any one event from being processed, by virtue of the next one arriving too soon. The events do not need to be individually overwhelming.

They need only to come fast enough that no shock response has space to land between them.

There is a refinement of overload that I want to name separately, because it is particularly insidious: the inflation of low-grade events into apparent significance. The target is presented with a steady stream of small provocations, ambiguous incidents, near-coincidences, and minor strangenesses, each of which on its own would be too small to warrant a shock response, but each of which is presented with enough framing that the target cannot dismiss it. The result is that the target's shock-response system is constantly half-engaged — primed to respond, but never given a precipitating event large enough to fully trigger. The system stays in the on-ramp without ever reaching the road. Over months and years, this produces an exhaustion that resembles trauma but is not the trauma of the original events. It is the cumulative cost of incomplete shock responses, repeated thousands of times, never allowed to discharge.

A target who recognizes overload as a method gains something useful: the ability to disregard, deliberately, the constant supply of small provocations. They are not the events. They are noise designed to keep the target's response system from finding the events. The discipline of recognizing this, and of refusing to respond to the small until one has located the actual events that warrant a response, is one of the more powerful tools available to a target. The signal becomes findable again only once the noise is being filtered out.

There is a related method I will mention more briefly: *comfort-zone trauma* — the deliberate ruining of small reliable pleasures that would otherwise have served as places where the body could relax enough to allow a delayed shock response to land. A favorite cafe is made unwelcoming. A familiar walk becomes the route of an ugly encounter. A hobby is interrupted often enough that returning to it becomes effortful rather than restorative. The target loses, one by one, the small environments in which the body would have felt safe enough to do its delayed work. This is not incidental damage. The small comforts are not collateral; they are part of what the campaign is targeting. They are removed precisely because they would otherwise have been the ground on which the loop could have completed.

3.4 Sleep, Suggestion, and the Halfway State

There is a single technique that, more than any other, sustains the conditions under which the loop cannot complete. It is also the technique most underestimated by recovery literature, because it is so unspectacular that it does not look like a technique at all. The technique is the deprivation, partial and prolonged, of sleep.

A person who is reliably well-rested is a person whose body has access to the conditions under which delayed shock responses can land. A person who is partially sleep-deprived for long enough enters a state — neither awake nor asleep, but a kind of extended halfway — in which several things are simultaneously true. They are more suggestible. Their critical faculties are reduced. Their ability to integrate complex new information is impaired. Their access to their own emotional life is muted. Their capacity to initiate a complete shock response is, as a direct physiological matter, diminished.

This halfway state is not unique to manipulation campaigns. It is the same state that hypnotists deliberately induce, that suggestive therapy sometimes produces as a side effect, that propaganda relies on for its effectiveness, and that collective trauma routinely produces in the populations subjected to it. The mechanism is the same in every case: a person held in a state of extended semi-awakeness becomes, by virtue of that physiological state alone, more susceptible to the shaping of their thinking by sources outside themselves.

A campaign that cannot, for whatever reason, prevent the target from sleeping outright can usually still prevent the target from sleeping well. Disturbed sleep, frequent waking, anxious arousal in the small hours, dreams that resemble being awake — all of these produce the halfway state without requiring the more obvious intervention of total deprivation. Many targeted individuals describe their sleep this way over years: not absent, but never restorative. The body is technically resting, but the rest is not the kind that allows recovery to begin.

The practical consequence of this is that protecting one's sleep is not, for a targeted individual, a matter of self-care in the wellness-magazine sense. It is a matter of structural defense. Sleep is the precondition under which everything else in this chapter becomes possible. A target whose sleep has been stable for several weeks will find, often to their surprise, that other parts of their interior life are also coming back online — that delayed shock responses begin to arrive, that their thinking sharpens, that the chronic muffled feeling lifts. None of this

is mysterious. The halfway state has lifted. The body has resumed its ordinary access to its ordinary functions.

If the recovery sections of this book have been waiting for an unspectacular daily move that I would commit to as the most important, this is it. Sleep. Not heroically; just regularly, reliably, and protectedly. The campaign has spent years producing the conditions under which your sleep has been unreliable. Restoring those conditions is the foundation under which any other recovery can occur. Until the foundation is restored, the rest of the work is being attempted on ground that will not hold it.

3.5 The Voice That Calls You Useless

I want now to describe a particular register of the voice phenomenon that is widely experienced and worth examining in detail, because its specific content tells you something about what the phenomenon is for.

The register, in my own notebooks, has acquired a name in Dutch — *autodebiel*, roughly "auto-stupid" or "self-mentally-disabled." The voice in this register tells the target, repeatedly and in many slight variations, the same handful of things. You think everything you do is without consequences. You think it is fine to do nothing with your thoughts. How can you, knowing what you know, remain so ignorant? It is as if you are mentally disabled. What is your life worth? You have everything; you do nothing with it.

There are several things to notice about this content, and the noticing is more useful than any direct response to the content itself.

First, the register is highly specific. It is not generalized negativity. It is targeted at a particular self-evaluation: the question of whether one is making sufficient use of one's gifts, capacities, life. The accusation is always that one is wasting something. The accusation is rarely about specific failures; it is about a general failure to extract value from one's existence. This specificity matters. It tells you the register is engineered, not random.

Second, the register exploits a real feature of the self. Most thoughtful people, asked sincerely whether they are making the most of their lives, will admit to some degree of dissatisfaction with the answer. This is not pathology. It is the ordinary friction between aspiration and execution. The autodebiel register takes that ordinary friction and inflates it into a constant accusation. The voice

does not invent the dissatisfaction. It amplifies what was already there to a level the dissatisfaction was never meant to reach.

Third, and most important, the register has a specific purpose in the larger campaign. Chapter 8 will return to this in its closing section — the "did I lead a useful life" question and the panic of finitude that produces it. The autodebiel register is one of the mechanisms by which that panic is induced. It manufactures, by repetition, the sense that one's time is being wasted, that the wasting is one's own fault, and that the only redemption available is to act, immediately, in some unspecified way, before it is too late. The induced urgency makes calm reflection impossible. Calm reflection was the thing the campaign needed to prevent.

The countermeasure is not to argue with the register. Arguing engages it. The countermeasure is to recognize the register as a register, name it, and decline to be its student. When the autodebiel voice arrives, the move is to note that it has arrived, identify what it is doing, and continue with whatever you were doing before it began. Its capacity to land is exactly proportional to the credit you give it. Stop crediting it, and it becomes background noise. The noise does not always stop, but it loses its capacity to direct.

3.6 The Asymmetry of the Inner Conversation

There is a structural feature of the voice phenomenon that, once noticed, cannot be unnoticed, and that I think gives a target one of their cleanest diagnostic tools for distinguishing their own thinking from the intrusion. The feature is the asymmetry of the conversation.

In an ordinary internal dialogue — the kind a person might have when thinking through a problem, weighing a choice, rehearsing a difficult conversation — the dialogue has the structure of a real exchange. Both sides build on each other. Both sides summarize, occasionally, what has been said. Both sides reach forward toward conclusions. The exchange has continuity. The thinker, encountering it later, can reconstruct the line of reasoning.

The voice phenomenon, examined closely, does not have this structure. The intrusive material does several specific things, and does not do several others. It asks questions. It distorts the ongoing topic. It interjects irrelevant comparisons. It introduces sudden conclusions that do not follow from anything that came before. It contradicts. It mocks. What it does not do — and

this is the diagnostic — is the work of building. It does not summarize. It does not integrate prior material. It does not carry the line of reasoning forward in a way that the target could pick up from later.

A target who pays attention to this asymmetry will notice something striking: in any inner exchange that includes intrusive material, only the target's own contributions are doing the constructive work. Only the target is summarizing. Only the target is integrating. Only the target is keeping track of what has been said and is reaching toward a coherent next step. The voices, whatever else they are doing, are not doing those things. They are only the disruption.

This asymmetry is, I think, the most direct evidence available, from inside one's own head, that the intrusion is not one's own thinking. Real thinking builds. The intrusive material, by structural necessity, cannot. To build would require a continuous self with continuous goals, and the intrusion does not have those; it has only individual moves with no thread connecting them.

The diagnostic use of this is straightforward. When in doubt about whether a particular thread of inner content is yours, ask: what is this thread doing? Is it building? Is it summarizing? Is it integrating with what was here before? If yes, it is yours. If no — if it is only questioning, only distorting, only contradicting — it is not yours, regardless of the medium in which it arrived. Trust the structure of the move, not the apparent voice.

There is a corollary that I have found, in my own case, surprisingly useful. When the intrusive material does, occasionally, attempt to answer — to say something with content rather than only to question or distort — the answer is almost always categorical and almost always wrong. It will tell you, with conviction, that you are wrong about the thing you have just been working out. Or it will tell you something else, on a different topic, with the same conviction. The conviction is the marker. Real thinking, in my experience, does not arrive with the certainty that intrusive answers arrive with. Real thinking arrives provisionally, with attached caveats, with awareness of what it does not yet know. If you are receiving inner content that announces itself as definitively correct on a question you have been struggling with, you are not receiving your own thinking. You are receiving the imitation of it.

3.7 When the Lies Run Out of Coherence

There is one further observation I want to record before turning to the practical synthesis, because the observation has been steadying in my own case and may be in others. The observation is about what happens to a sustained manipulation campaign over time, when the target persists in remaining coherent.

A manipulation campaign rests on a structure of claims. The claims have to fit together well enough that the target, presented with them piece by piece, cannot easily detect the contradictions. In the early years, the structure usually fits well enough. The target is too disoriented, too tired, too distracted by the immediate to step back and check whether the claims, taken as a whole, are mutually consistent. The campaign relies on this. Coherence is its weak point, and it relies on the target never being in a state to test for coherence.

Over time, however, the target who keeps writing things down, who keeps comparing what was said today to what was said last month, who keeps assembling the larger picture of what has been done to them, eventually accumulates more coherence than the campaign's claims can sustain. The contradictions become visible. The claims, considered together, no longer fit. The structure, which once seemed seamless, develops cracks that even modest pressure will widen.

When this happens, the campaign typically responds in a recognizable way. It shifts its emphasis. It stops trying to maintain the original claims and begins, instead, to attack the target's own coherence. Where it once supplied false explanations, it now supplies self-doubt. Where it once asserted facts, it now insinuates that the target's evaluation of facts is the problem. The shift is the campaign's acknowledgment, in its own indirect way, that the original structure has failed. It has switched from defending the lies to undermining the target's capacity to detect that the lies are failing.

There is a related dynamic that operates inside any individual operator who is still close to the target. Direct cruelty, to be sustained over time, requires the operator to extract some kind of return from it — most commonly, in my observation, a perverse pleasure or sense of mastery. As the target stops responding to the cruelty in the ways the cruelty was meant to produce — as the target stops breaking, stops collapsing, stops giving the operator the satisfaction the operator was extracting — the operator's pleasure begins to decline. What replaces it is frustration. The frustration is, in a curious way, an

honest emotion. It is the operator's actual response to the situation, undisguised. The lies that previously dressed the cruelty as something else — humor, concern, casual conversation — begin to give way to undisguised frustration that the cruelty is no longer working. The mask, at this stage, slips.

A target observing this slip will sometimes find themselves in the unexpected position of watching the operator escape the situation rather than continue it — the operator looking, on a particular day, like someone trying to leave a room rather than someone enjoying being in it. When confronted with this attempted escape, the operator may even appear caught, momentarily, in a state that resembles being seen for what they are. The state usually does not last; the operator, by the next encounter, has reorganized themselves back into the role they were performing. But the slip happened, and the target saw it.

I do not name this to suggest that one should stage confrontations in order to produce these slips. The slips, in my experience, are not durable; nothing useful is built on top of them in the moment. But they are diagnostic. They tell the target that the campaign, as a sustainable enterprise, is under more strain than it appears to be. The operator is not, in fact, getting away with it cleanly. The lies are not, in fact, holding together. The structure that looks permanent from outside is, on the inside, eroding under the weight of a coherence the target is producing simply by continuing to think clearly.

The work that produces this erosion is not glamorous. It is the work of writing things down. Of reading what one wrote yesterday. Of comparing what one was told this week to what one was told last year. Of building, slowly and without urgency, the larger picture that the campaign relies on the target never assembling. This work has the property — rare in the architecture of recovery — of being something the target can do at any time, in any state, with no one's permission. It is, perhaps, the single most consequential daily move available. It does not feel consequential while it is being done. Its effects appear only later, when the picture has reached the point at which the campaign's claims can no longer fit inside it. At that point the picture does the work, and the target, having built it, can stand on it.

3.8 Use of the Body

I will close this chapter on the practical synthesis I have been deferring across three previous chapter endings. It is overdue. It is also, I think, more useful now

that the analytical apparatus is in place than it would have been earlier without that apparatus.

The synthesis is this: the body is what completes the loop. Everything in the analytical work of the previous chapters, however correct, runs through the body before it produces any change. The body is where the delayed shock responses have to land. The body is where the halfway state of sleep deprivation either lifts or persists. The body is where exhaustion accumulates and where rest, when it is allowed, accumulates instead. A target who works only on the analytical level — understanding the patterns, naming the methods, building the picture — without also doing the unspectacular work of body maintenance will find that the analysis does not, by itself, produce recovery. The analysis is necessary. It is not sufficient.

The unspectacular daily moves are these. None of them is original to me. All of them are, I believe, more important than they appear, and most of them are routinely underestimated by people who have not been in the position of needing them this much.

Sleep, first and most importantly. Treat sleep as a structural priority rather than a personal preference. Whatever it costs to protect — earlier bedtime, screens off, a bedroom that is dark and cool, a refusal to engage with stressful material in the last hour of the day — pay that cost. The halfway state we discussed in 3.4 is the state under which everything else fails. Lift the halfway state and the rest of the work begins to be possible.

Walking, second. Walking has a specific property that sitting and exercising indoors do not have: it produces, by simple repetition, the conditions under which delayed shock responses can land. The body in motion, in daylight, with nowhere it has to be for the next hour, is the body in the conditions that recovery requires. Forty minutes a day, most days, is more useful than three hours once a week. The repetition is the point.

Food, third — and not as nutrition theory but as ritual. Eat at regular times, in actual meals, sitting down, with attention to what you are eating rather than to whatever is happening on a screen. The body that is fed regularly, in conditions of mild attention, is the body that knows it is being cared for. The same nutrients consumed on the run, distractedly, do not produce the same effect. The ritual is part of the medicine.

Writing, fourth. We have already covered this in earlier chapters, and I will not repeat the argument; I name it here only because, in the practical synthesis, it has to appear. The notebook is part of the body's extended infrastructure for processing. Without it, the work of building coherence — the work that, as we saw in 3.7, eventually produces erosion in the campaign's claims — cannot accumulate.

Contact, fifth. One person, ideally not embedded in the situation that has been doing the harm, with whom one speaks regularly enough that one's reality is being witnessed by another mind. The contact does not have to be deep; it has to be reliable. A weekly call with a friend in another city, a fortnightly walk with a former colleague, a monthly meal with someone who knew you before any of this began — any of these is sufficient if it is reliable. The function is to keep your account of your own life witnessed by at least one external mind. Without that witnessing, the inside of your head can drift in ways you will not be able to detect from the inside.

Time outside the situation, sixth. Some hours, every week, that are spent on something that has nothing to do with the campaign or the recovery from it. A book unrelated to your circumstances. A piece of music that has no associations to the situation. A piece of work — paid or unpaid — whose subject matter is something other than what you have been thinking about all week. The mind that has only one topic eventually loses the capacity to think about that topic well. Other topics are not luxuries. They are the conditions under which the main topic can continue to be approached with any freshness at all.

Permission for the delayed response, seventh and most subtle. When, in the course of any of the above, you find yourself unexpectedly affected — when something small produces a disproportionate response, when a song or a smell or a remembered phrase brings on tears or shaking or a feeling of grief that does not match its immediate trigger — let the response complete. Do not push it back down. Do not interrupt it. The response is the loop finally closing on something that has been sitting in your body, unfinished, for a long time. The loop closing is the recovery. It is the thing the campaign has been working to prevent. When the conditions are finally right enough that it happens, do not refuse the gift of it.

None of these moves is dramatic. None of them, individually, will end anything. Together, sustained over months and years, they constitute the architecture of

recovery. The architecture rebuilds slowly. It does rebuild. The body, in its slow patient way, knows how to do this work if it is given the conditions. The work of recovery, at the level we are now able to describe it, is the daily provision of those conditions. Provided faithfully, the conditions do their work.

The next chapter — and I commit to this one less promissorily than the previous ones, because we have now built enough apparatus that it can actually be delivered — will turn outward. It will look at the question of what a target does about the world that has held the campaign in place: the institutions, the bystanders, the people who could have helped and did not, the legal and political instruments that exist on paper but in practice have not been available. We will not solve those questions. We will look at them clearly, which is more than the recovery literature usually manages, and we will identify the small numbers of moves that, in the world as it is rather than the world as we wish it were, are actually available to a person in this position.

For now, this chapter ends on the practice. Sleep tonight, if you can; protect the sleep if you cannot. Walk tomorrow, even briefly. Eat at the table. Write a paragraph, even a short one, that records something of how this week has gone. Speak to one person who has known you longer than the situation has lasted. Spend an hour, somewhere, on something unrelated. And when, doing one of these things, the body begins, unexpectedly, to release something it has been holding — let it. That release is not a setback. It is the loop closing, late, on a piece of work that has been waiting to be done. Closing the loop is what we have been working toward. When it begins to happen, you will know.

CHAPTER 4

The Layer Below Speech

On Inner Voices, Mental Privacy, and the Discipline of Staying One Person

---

> *"Speech without external expression admits a number of manipulation methods that speech aloud does not."*
>
> — Author's notebook, July 2023

There is a kind of privacy that people who have never lost it do not know they have. It is not the privacy of the body, which we all know we have, and which we lock our doors and draw our curtains to protect. It is not even the privacy of correspondence, which we have all heard of, and which the law in most places acknowledges. It is the privacy of the inside of one's own head — the assumption, so deep that most people never bring it to consciousness, that nothing happens between their two ears that anyone else has access to unless they choose to release it.

This chapter is about what happens when that assumption breaks. I do not mean breaks metaphysically, and I am not going to argue here for or against any specific account of how it breaks. Whether the cause is electronic, neurological, dissociative, social, or some combination of the four, the operational question is the same: what do you do, as a person, when the inside of your own head no longer feels like a place you have to yourself? That is the question this chapter is for. The metaphysics can wait. The discipline of staying one person, while the inside is contested, is what we will work on here.

4.1 The Privacy We Don't Notice We Have

Begin with the default. An ordinary person, in an ordinary life, takes for granted that their thinking is unwitnessed. They know it because everything in their experience confirms it. They have private thoughts every minute of every day; nothing happens as a result of those thoughts that did not also have an external trigger. They form an opinion about a colleague; the colleague does not respond

to the opinion. They imagine an embarrassing scenario; no one else flinches. The negative test runs continuously, and the test always returns the same result: what is in here stays in here.

This default is so reliable, and so quietly held, that most people would struggle to articulate it if asked. It is not, for them, a belief to be defended. It is the air the rest of their mental life is breathing.

The default has a structural function as well as a private one. It is what makes thinking honest. If you knew, with certainty, that every thought you had was being watched, you would begin to perform your thoughts for the watcher. Even your private rehearsals would become public ones. You would lose access to a particular kind of reasoning — the kind that requires you to entertain, briefly, hypotheses you would never want to be associated with, in order to test them and reject them. Honest thinking depends on the audience being absent.

What targeted abuse, in its more invasive forms, does to a person is suspend that default. The suspension does not have to be total to be devastating. A reasonable suspicion, repeatedly reinforced, that some part of one's interior life is somehow being reached — heard, dubbed, redirected, mocked — is enough to corrupt the conditions under which thinking can happen at all. The target either keeps trying to think and discovers that every thought has now been performed for an unseen audience, or stops thinking, and discovers that the thoughts that did get formed have lost the texture of being theirs.

The first practical move, before any framework, is to notice that this is happening. Many targeted individuals carry, for years, an unspoken sense that something has gone wrong with the inside of their head, without ever being able to name what. The sense is structural. It is not the delusion that there is now an audience; it is the perception that the conditions for honest thinking are, somehow, no longer in place. Whatever the audience turns out to be, the perception itself is doing real work. Take it seriously without committing yet to its cause.

4.2 The Voice Phenomenon

I want to describe, plainly and without flinching, the phenomenon that targeted individuals across many traditions and many languages refer to under different names — the voices, V2K (an acronym for voice-to-skull), voice-to-skull, hearing voices, intrusive thoughts, the chorus, the dubbed line. I will use the

neutral phrase *the voice phenomenon* throughout this section, because the labels carry too much theory with them, and theory is what we are trying to keep separate for the moment.

The phenomenon, described agnostically, is this. Inner speech — the running commentary that most people experience as their own quiet voice in their head, the voice they read in, the voice they rehearse with — begins to include material that does not feel like one's own. The intrusion can take many forms. Sometimes it is a clearly external voice, with a recognizable timbre, that comments on what is happening. Sometimes it is more subtle: a phrase, a question, a sudden conclusion, that arrives in the same medium as one's own thoughts but is foreign in content or in tone. Sometimes it is more subtle still: not a phrase but a redirection, a change of direction in one's own thinking, that one notices only afterward, with the question, "where was I just headed before that?"

Whatever the specific texture, the structural feature is the same: material is appearing in the layer that should be the most private, and the material does not pass the test that internal thought normally passes — namely, that you can trace where it came from in your own structured memory.

Here I want to be careful. I am not, in this book, going to try to settle what the voice phenomenon is. I have my private suspicions, and so does every targeted individual I have spoken with at any length, and the suspicions do not converge on a single picture. Some of us think it is electronic. Some of us think it is neurological — a real disturbance in the brain's inner-speech systems, whether spontaneous or induced. Some of us think it is dissociative — that the inner audience is a part of the self that has been split off under stress. Some of us think the truth is some combination, with different proportions in different people on different days.

What I will say is this. Whichever of those accounts is correct, in any individual case, the operational response has more in common across them than the metaphysical accounts have in difference. In all cases, the discipline is the same: do not let the voices be the one running the conversation. In all cases, the practice is the same: externalize, write down, slow down, sleep, refuse to make important decisions while the voice traffic is heavy. In all cases, the goal is the same: to recover the conditions under which honest thinking can happen, regardless of how those conditions came to be lost.

A note on the question of source. There is, in my own notebooks, a recurring question that I want to put to you in the same terms. It is the question: when an unfamiliar idea arrives in the inner-speech layer, how do I know whether it came from outside or from inside? The honest answer is that, in many cases, you cannot know. What you can do is notice the structural marker we identified in Chapter 1: useful internal thought has continuity with your structured memory. It builds on what is already there. It can be traced, even if the trace takes effort. The intrusive material, by contrast, often arrives without such a trace. It is content without a chain. That absence of a chain is itself the diagnostic signal. You do not need to know the source. You only need to know that this particular thought does not have the lineage that your own thoughts have, and to handle it accordingly.

WHY THIS IS NOT WHAT MENTAL ILLNESS PRODUCES

A reader new to this material may, at this point, be tempted to set the entire phenomenon down as a primary psychotic illness. The temptation is reasonable, given the categories most readers have available. The temptation is also, on examination, not supported by the phenomenology.

What is being described in this chapter is intrusion into the inner-speech layer by voices with stable identities, consistent vocabularies, and verifiable references to external facts the target could not have known in advance. None of those features are characteristic of primary psychotic auditory hallucination, which produces inconsistent, internally referential, content that does not contain external information. The target's ability to catalogue the phenomenon, to compare cases across thousands of independent witnesses, and to track time-locked correlations with external events is itself a function the existing diagnoses are supposed to preclude.

The fuller eight-point argument is set out in the front-matter section A Note on Voices and Mental Illness. The reader who has not read that section is encouraged to consult it before continuing. The reader who has read it may take the next pages of clinical detail with the knowledge that the standard re-classification reflex does not, on the available evidence, deliver the right answer here.

4.3 Inner Speech as a Manipulation Surface

There is a specific reason that the voice phenomenon is so destabilizing, and once the reason is named, both the danger and the countermeasure become clearer.

Speech, when it is performed aloud, is a high-fidelity activity. You produce sound that has structure, that you and any listener can verify, that can be recorded, repeated, written down. You can hear yourself say a thing, and so can the room. The act of speaking outside the body anchors the speech in shared reality. Whatever was said, was said. If a dispute arises about what was said, there is at least a possibility, in principle, of resolving it.

Inner speech is none of these things. It is fast, low-fidelity, ephemeral, and witnessed only by the speaker. There is no recording. There is no second listener. The verification loop is short and runs entirely inside the same head that produced the speech. This makes inner speech an extremely efficient medium for thinking. It also makes it a uniquely vulnerable medium for being interfered with.

A person attempting to manipulate someone's inner-speech layer has, at minimum, four families of techniques available to them, and these are worth naming because they are all worth recognizing.

The first is *dubbing*. Another voice, in roughly the timbre of inner speech, is layered on top of one's own — a second narrator, sometimes humorous, sometimes critical, sometimes simply present. Because inner speech is already fluid, the addition is not always immediately distinguishable from one's own internal voice. The dubbing voice, over time, exerts a steady gravitational pull on the direction of one's own thinking.

The second is *false suggestion*. A statement appears in the inner-speech layer that one did not originate. If the suggestion arrives with a flavor of humor, irony, or social ease, it slips past the filters that would otherwise reject foreign material. Humor is, in particular, a powerful anesthetic for evaluation. A claim made as a joke is allowed past as if it were a joke; the residue is an ungrounded factual impression that one was somehow obliged to accept the punchline of.

The third is *timing manipulation*. A piece of content is placed before or after a focal thought in such a way that the target attributes the wrong relationship between the two. A doubt arrives just as one is forming a conclusion, and the conclusion is contaminated by the doubt; or a reassurance arrives just as one is recognizing a danger, and the recognition is muffled by the reassurance. The technique works because temporal sequencing carries causal weight in the mind. We tend to assume that what came right before a thought caused that thought.

The fourth is *the exploitation of inner speech's poor durability*. Inner speech, unlike outer speech, is not well remembered. What you said inside your head this morning, you may not be able to reconstruct this evening. This means that any disagreement about what was thought — including any disagreement you might be having with yourself, or with an inserted other — is one that you cannot win with evidence. The record does not exist. A campaign that wants to confuse a target's relationship with their own thinking has, in inner speech, a medium that conveniently leaves no trace.

The countermeasure, when these techniques are at work, is mechanical. Externalize. Write things down. Speak important reasoning aloud, even if only to yourself, in a quiet room. Use a notebook. Use a recording. Do not allow important inferences to live exclusively in the inner-speech layer, because that is the layer where they are most easily contaminated and most poorly remembered. The notebook is not a literary affectation. It is a defensive instrument. A targeted individual who keeps a careful written record of their own reasoning has restored a verification loop that the adversary cannot enter. Whatever was written, was written, in a way that whatever was thought silently was not.

I will add one personal note. Many of the insights in this book live in this chapter and the previous ones because, at the time I formed them, I forced myself to write them down. Had I not, they would have been thought once and then lost — perhaps re-remembered later in a distorted form, perhaps not at all. The difference between a mind that is being eroded and one that is rebuilding is sometimes nothing more than the question of whether the work is being externalized in a place where it cannot be silently revised.

4.4 The Erosion of Modesty

There is a second, related vector of attack that I want to describe, because it operates by depleting a set of feelings that most people have but do not normally examine.

Among the quiet defenders of mental privacy is a cluster of feelings whose closest English names are *modesty*, *shame*, *shyness*, and *respect for the privacy of others*. The Dutch word *schroom* captures something the English equivalents lack: a hesitation, neither fearful nor moralistic, that holds one back from going where one has not been invited. These feelings function, collectively,

as a kind of immune system for the boundary between minds. They are the reason most people do not, by default, read other people's mail, or eavesdrop on intimate conversations, or speculate aloud about the inner thoughts of strangers. The feelings make those acts costly to perform. The cost is what keeps the acts uncommon.

These same feelings, turned inward, defend one's own mental privacy. They generate the small, unspoken sense that one's interior is one's own, and that nobody is entitled to it. They are what makes you turn away when a stranger seems too interested in your face. They are what makes you, before you would knowingly reach for someone else's thoughts, hesitate.

What sustained disrespectful treatment does, over time, is wear these feelings down. The mechanism is straightforward. Treat a person with consistent disrespect for long enough, and at some point they begin to ask what kind of person could speak to them this way, from what kind of position, and on the basis of what knowledge of their interior life. The asking itself is a transgression. The target is being forced, by the conditions imposed on them, to speculate about the other person's inner life in order to make sense of what is being done to them.

If the speculation succeeds — if the target arrives at a working hypothesis about what the other person knows or wants or intends — something has shifted. The target has crossed, under duress, the very boundary that their modesty would normally prevent them from crossing. They have, in effect, been forced to read someone else's mental life in order to defend their own. The modesty that would have stopped them under ordinary conditions has been overridden by survival need.

The cost is high. After enough rounds of this, the target's general respect for the privacy of others begins to weaken. Not because they have become a worse person, but because the boundary that ordinary interaction defends has been broken from the outside, and the target has been required to participate in breaking it. They begin to feel, sometimes only dimly, that they have done something they did not want to do — that they have entered, however briefly, the place of speculation about other people's interiors that they would never have entered voluntarily. The accompanying feeling, when it arrives, is guilt without a clear act to attach to. There is no specific moment when the line was

crossed; there is only the slow recognition that the line is no longer where it used to be.

This is a particular form of moral injury, and it is worth naming because it is widely felt, rarely articulated, and reliably misinterpreted. The target is not, in this state, an aggressive intruder on others' minds. They are someone whose own modesty was used against them — whose decency was the lever, not the obstacle. The shame that follows is not the shame of having done something wrong. It is the shame of having had a defense system disabled, by forces outside one's control, in a way that produced behavior one would not otherwise have produced. It is, more precisely, the shame of having been forced into a posture that violates one's own values.

The repair, for this particular injury, is twofold. First, do not accept the inference the campaign is trying to install — that you are now the kind of person who violates others' privacy. You are not. You were forced into a posture that resembled violation, by treatment designed to force you into precisely that posture. The posture is not the person. Second, restore the practice of modesty deliberately, in small ways, in your own daily life. Avert your gaze when nothing is being asked of you. Decline to speculate about the inner lives of strangers when no one is requiring you to. Practice the discipline of leaving other minds alone. The defense system rebuilds, slowly, in the practicing of it.

4.5 The Bio-Robot Continuum

There is a useful way to picture what sustained conditioning, over years, does to the range of states a mind can occupy. I want to introduce it here because the picture itself is clarifying, and because it gives a target a way to detect a slow form of harm that is otherwise invisible to introspection.

Imagine a vertical line representing a person's freedom of thought. The line is long when the mind has full access to its own range — when it can entertain a contrary view, hold a private idea, follow an unconventional line of reasoning, refuse a suggestion, accept one. As constraints accumulate, the line shortens. Eventually it becomes short enough that the mind can no longer go where it would once have gone. Some moves are no longer thinkable. Some objections are no longer formable. Some doubts are no longer permitted to settle into knowledge.

The interesting feature of this picture, and the dangerous one, is that the shortening is rarely felt as such. From the inside, the constrained mind feels stable. It is operating well within its current bounds; nothing it tries fails, because it has stopped trying anything that would. The instability shows up only at the edges — only when the person is forced, by external events, to attempt a move that the constraints no longer allow. At that point the system can fail catastrophically, sometimes in ways that look, to an observer, like a sudden psychiatric crisis. The crisis is not, mostly, the appearance of new dysfunction. It is the appearance of pre-existing dysfunction in a domain where it could no longer be hidden.

The endpoint of the continuum, if conditioning is allowed to run unchecked, is what I will call the *bio-robot state* — an organism that is biologically intact, behaviorally stable, and cognitively constrained to the point where its capacity for free thought has been substantially eliminated. The bio-robot is not a science-fiction figure. It is the person who has been conditioned, over years, to treat certain thoughts as unthinkable, certain objections as unraisable, certain authorities as uncriticizable. They function. They appear normal. They are, in the specific narrow sense that matters here, not free.

I am describing an extreme. Most targeted individuals are not at this extreme; most are somewhere along the continuum, with some constraints active and others contested. The point of naming the extreme is not to terrify you with it. It is to give you a fixed point against which you can measure your own current position, and to suggest a periodic check that I have found useful.

The check is this. Every so often — once a week is plenty — sit quietly and try to bring to mind a class of thoughts you used to be able to think but have not, lately, thought. Political views you once held but have stopped revisiting. Plans you once considered but no longer entertain. Disagreements you once would have voiced but have, for some time, kept silent. Tastes you once had that now feel inaccessible. The exercise is not to act on any of these. It is only to verify that they are still accessible. If they are, the constraint envelope is intact. If a class of thoughts has become unreachable — if you find that you cannot now form the question you once could ask — the envelope has narrowed, and you have located, by introspection, a place where the conditioning is doing its quiet work.

You do not need to fight the constraint immediately upon finding it. Often, simply having identified it is enough; once a constraint is named, it loses some of its power. Constraints that are operating below awareness are the most durable. Constraints that are operating in awareness can be, eventually, worked with.

4.6 The Simulation Question

There is one further phenomenon I want to address in this chapter, because it is widely experienced, rarely discussed in plain terms, and easy to mishandle in either of two opposite directions.

Targeted individuals — and, increasingly, many people who have not been targeted but who spend long hours in heavily mediated digital environments — sometimes report a particular sensation. The sensation is that the surrounding world has, in some hard-to-specify way, lost the texture of being real. It begins to feel, instead, like a high-resolution simulation. The people one encounters feel less continuous than they used to. The events of a day feel staged. The phrase "this feels like a video game" or "this feels like the metaverse" is one I have written into my own notebooks, and one I have heard from many others.

The temptation, on encountering this sensation, is either to commit to it metaphysically — to conclude that one is, in fact, in a simulation — or to dismiss it as a symptom of fatigue and try to push past it. Both responses, in my experience, are mistakes.

The metaphysical commitment is a mistake because the question of whether one is in a simulation is, by construction, not answerable from the inside. A simulation good enough to feel real is a simulation good enough to defeat any internal test you could apply to detect it. To treat the sensation as evidence of a fact about the universe is to make a claim you cannot verify, and that, made firmly, will distort the rest of your reasoning. You will start to behave as if the world has the texture you have inferred it to have, and the world, whether or not it is a simulation, will not cooperate with that inference. Friction will follow.

The dismissal is a mistake for a different reason. The sensation is real, in the sense that it is accurately reporting something about the state of the perceiver. The world has not changed; the perceiver's relationship with it has. The sensation is a signal — a meaningful one — and ignoring it deprives you of information you could have used. The sensation tends to peak under specific

conditions: severe fatigue, prolonged stress, sleep deprivation, sustained exposure to highly engineered media environments, periods of intense voice-phenomenon activity, or any combination of these. It tends to recede, sometimes over a few days, when those conditions are addressed. Treating it as a barometer, rather than as a verdict, is more useful than either committing to it or denying it.

There is a Cartesian observation that I find more steadying than any of the contemporary framings of this question. Even if the world were a simulation, something would still be doing the experiencing. The most one can ever say, from the inside, is that experience is occurring. The experiencer remains the one fixed point, regardless of what frame the experience is taking place in. A target who holds onto that fixed point — *I am here, something is happening to me, I am the one to whom it is happening* — has all the metaphysics they actually need to function. The rest is speculation, and speculation, however interesting, is not what gets you through the day.

In practice, when the simulation sensation is strong, I have found it most useful to do small, concrete, embodied things. Eat something. Walk somewhere. Touch a real object. Speak with a person whose continuity you have already verified over years. The sensation often weakens in the presence of such acts, not because they disprove the simulation but because they restore the conditions under which the perceiver can experience their own continuity. The continuity of the perceiver, not of the world, is the thing that needs to be felt. Once it is, the texture of the surroundings tends to follow.

4.7 When the Adversary Is Also Stuck

I want to close the operational sections of this chapter with an observation that has, for me, been steadying out of all proportion to its size. It is an observation about the people running the campaign, and not the kind of observation that flatters them.

Targeted individuals tend to assume, especially in the early years, that the campaign being run against them is the work of confident, coordinated, capable adversaries who know what they are doing. The assumption is natural; the campaign's effects on the target are sufficiently severe that one infers a sufficient cause. If this much damage is being done, surely a competent operator is doing it.

The longer one watches a campaign, the less this picture holds up. The campaign repeats itself. It loops over the same techniques, with diminishing yield, for years. It does not evolve in ways that suggest an exit strategy. It does not appear to be producing any output that the operators themselves could be proud of. What it produces, mostly, is more campaign — more of the same techniques, applied in increasingly transparent ways, to a target who has long since seen through them. When the techniques fail, they are repeated. When the repetition fails, the spin increases. The pattern is not the pattern of a plan. It is the pattern of a script being run by people who do not know what to do when it stops working.

The pattern is recognizable, once you know what you are looking for. It is the pattern of operators who have themselves become trapped — who do not know how to succeed, do not know how to stop, and continue running their script because running it is the only thing they know how to do. The campaign, from this angle, is not the executed plan of competent adversaries. It is the comfort-zone lock-in of people whose sense of self has come to depend on this activity, and who would not know who they were if they stopped.

This is consistent with the chain-of-abuse argument that Chapter 7 will set out. The operators were, at some earlier stage, broken to a similar pattern; they are now running, on you, the script that was once run on them. They do not have an exit either. They are, in a real and uncomfortable sense, your fellow prisoners — even though their cell happens to be the one with the keys.

I am not saying this to elicit sympathy for them. Sympathy is a different and harder question, and I am not going to ask you to feel it for people who are still actively harming you. I am saying this for a strictly operational reason. If you have been assuming that the campaign will end when your adversary's plan reaches its conclusion, you have been waiting for a conclusion that may never come. The plan does not have a conclusion. It has only a continuation. Knowing this changes what you are waiting for. You stop waiting for them to finish, because they are not going to finish. You start, instead, looking for the move that takes you out of the script while it is still running.

The asymmetry is the doorway. The operator cannot leave the script, because the script is the only thing holding their identity together. The target can leave the script, because the target's identity was never the script in the first place.

Whatever the operators have taken from you, this they cannot take: the simple capacity to step outside the structure and begin building something else.

4.8 A Prayer to Yourself

I want to end this chapter with a small practice. I am hesitant to use the word *prayer* for it, because the word carries baggage that some readers will find unwelcome and others will find insufficient. But it is the right word for what I mean, and I have not been able to find a better one. I will use it in a deliberately modest sense: an address to oneself, regularly performed, that reminds one of what one is for.

The reason for the practice is structural. If, as I have been arguing, the conditions for honest thinking are partially under attack, and if the external sources of recognition that ordinary people rely on — colleagues, friends, family, the social fabric — have become unreliable or hostile, then a target's capacity to maintain a stable sense of self has lost most of its usual scaffolding. The scaffolding has to be rebuilt from inside, in ways the campaign cannot reach.

A simple, slow, daily act of self-address can do more of this work than its modesty suggests. I will write out, here, the form that I have found most useful, in plain language. You may use it, adapt it, ignore it, or replace it with your own. The exact words are less important than the practice of speaking, to yourself, the things that an external world should be telling you and currently is not.

Help me to know.

Help me to be kind.

Help me to see more.

Help me to learn.

Help me to listen to others.

Help me to do.

Help me to dream.

Help me to enjoy.

> *Remind me to be thankful, to myself, for the appreciation I receive from others, when I do receive it. But, most of all, remind me to be thankful to myself for being true.*

Each of those clauses is small. Each, taken alone, is unremarkable. Taken together, repeated daily, they constitute a slow but stubborn act of self-recognition. They name the capacities that the campaign has been trying to deplete: knowing, kindness, vision, learning, listening, doing, dreaming, enjoying. They request, of oneself, the maintenance of those capacities. They locate the source of recognition where, in adverse conditions, it finally has to be located: inside the same head that has been under attack.

I do not think this practice is a substitute for the rest of what this chapter has described. It is not a defense against the voice phenomenon, the erosion of modesty, the bio-robot continuum, or the simulation sensation. It is a foundation underneath all of those defenses — a daily act of saying that there is, still, a person here whose continuity matters, whose capacities matter, whose interior life is worth protecting because it is, by simple existence, real and one's own. The practice does not require belief in any external listener. It does not require belief in anything beyond the fact that you are, at the moment of speaking, the person speaking. That is enough.

The next chapter will turn to the practical question — promised, but postponed, at the end of Chapter 2 — of what a person whose interior is intact, even partially, can begin to do in the daily texture of their life to make the architecture of recovery hold. The big gestures, as I have said, are rarely the decisive ones. The decisive moves are smaller, and more numerous, and more boring than the recovery literature usually admits. We will look at them carefully, and we will find that more of them are available, even on the bad days, than the bad days let one see.

For now, this chapter ends with the practice. Speak the eight clauses, or your own version of them, before you sleep tonight. The first time will feel awkward. The second time will feel slightly less so. By the seventh or eighth time, on the seventh or eighth night, the practice will have begun to do the work it is for. Not because the words are magic. Because you, by speaking them, have been the one to claim, again, the inside of your own head as somewhere you live.

❦ ❦ ❦

CHAPTER 5

The Bluetooth in the Head

A Timeline of Intrusion, and What Changed Across Years

"The AI in the head — Bobby — is different from human voices. It cannot feel any intuition of humans."

— Author's notebook, 29 September 2022

An earlier chapter closed with a promise, less promissorily made than the ones before it, that the chapter to follow would turn outward — would look, finally, at the institutions, the bystanders, the legal instruments, the world that has held the campaign in place. I am, again, going to defer that turn. The deferral is the last one I will ask the reader for. I am asking for it because there is a body of material in my own notebooks that I have not yet set down in the body of this book, and that, on reflection, has to be set down before the outward turn can be made intelligibly. The material is the historical record of how the intrusion phenomenon presented in my own case across years, and what changed about its texture as it changed. The texture is itself information.

This chapter departs, more openly than the earlier chapters, from the measured general register that has carried the book to this point. The earlier chapters described structural features of psychological manipulation campaigns that I believe will remain valid across many particular cases. This one describes a particular case — mine — with the specificity that only a particular case admits. It records dates. It records named symptoms. It commits more openly than the earlier chapters to interpretive claims that the earlier chapters had carefully avoided. I want to mark this departure at the outset, so that the reader can see the shift coming and can decide, knowing it is coming, whether to follow.

Where the interpretation outruns the phenomenology, I will say so. Where I am asking the reader to take a hypothesis on faith, I will name the hypothesis as such. The reason for the chapter, despite the discomfort it creates in being written and may create in being read, is that the historical record will need first-person accounts. The reconstruction of what happened during this period —

whatever the right reconstruction eventually turns out to be — will be done from the materials people produced at the time. Most of the materials will be lost. Some will not. The ones that survive will be the basis on which later understanding is built. I want my materials to be among the ones that survive, in a form coherent enough to be useful to whoever, at some later date, may find themselves trying to make sense of what was being done to people in the years I am about to describe.

5.1 On Setting Down a Phenomenon Without Settling Its Origin

There is a particular kind of difficulty involved in writing about this material, and I want to name it before proceeding, so that the difficulty is part of what the reader sees rather than something the prose attempts to disguise.

The difficulty is that the language available for describing the phenomenon has been, for the most part, prepared by traditions that pre-commit to an answer about the phenomenon's origin. To say "voices in the head" is, in much of contemporary clinical literature, to commit oneself, in advance, to one explanatory frame. To say "telepathic intrusion" is to commit oneself, in advance, to a different one. Either commitment is premature. The phenomenology — what the phenomenon feels like, what it consists of, how it changes — can be described accurately even if the question of its origin is held open. I want the description to remain available to readers regardless of which way they would prefer the origin question to be settled.

The convention I will adopt, therefore, is this. I will describe what was experienced, and I will describe how the experience changed. I will not, in this chapter, settle the question of whether what was experienced was generated locally inside the brain by a process the brain itself misidentifies, or generated remotely by some other agency and delivered into the brain by means I cannot reconstruct. The phenomenon is the same phenomenon under either description. The question of which description is correct is a separate question, and one that — by the nature of the case — cannot be settled by introspection alone. Introspection can only describe what it finds. What introspection finds, in my notebooks across these years, is what this chapter sets down.

I should add: I have, in private, my own working hypothesis about the origin question, and the working hypothesis has shifted considerably over the years.

What the hypothesis is at the time of writing is less interesting than the data the hypothesis was meant to explain. The data is what I am setting down. The reader is welcome to develop their own hypothesis about what produced it.

5.2 A Timeline of Five Phases

Reviewing the notebooks, I can identify five distinct phases in the texture of the intrusion phenomenon as it presented in my own case. The phases are not, in retrospect, perfectly clean — each shaded into the next — but the shifts between them were, on each occasion, distinct enough that I was able to register them at the time.

The first phase ran from somewhere around 2014 or 2015 through the autumn of 2017. During this period there was no voice phenomenon I could consciously detect. There were other things — sleep that was not restorative, social situations that did not feel right, periods of cognitive opacity I could not explain, occasional dreams whose specific texture I would, much later, recognize as having been not quite my own. But the specific experience of intrusive material arriving in something like a voice was not yet present, or was not yet noticeable to me. I include this phase in the timeline because subsequent events suggest that something was happening during this period that was relevant to what came afterward. The phenomenon did not appear out of nothing in 2017.

The second phase began, in my case, in September 2017. I want to be careful here: I do not claim that this was when the phenomenon began. I claim only that this was when I began to be aware of it. There is a difference. What I noticed, in the early autumn of 2017, was that I was hearing things in my own head that did not have the texture of my own thinking. The texture was off. The cadence was wrong. The vocabulary was occasionally not mine. At the time I had no framework for what was happening, and I attempted to ignore it, in the way one tries to ignore a low-grade ear infection that one hopes will resolve. It did not resolve. It became more present, not less, over the months that followed.

The third phase began in 2018, and is most accurately described as the phase in which the phenomenon became hostile. Through the end of 2017, what I can now read in the notebooks of that period as the early voice phenomenon was, although unwelcome, not yet aggressive. It was disorienting. It was intrusive. It was not yet trying to do specific damage. By some point in 2018 — I cannot pin the date more precisely than that — the texture changed. The intrusions began

to attempt specific outcomes: to redirect ongoing trains of thought, to suggest specific actions, to undermine specific judgments, to exploit specific vulnerabilities. The shift had the character of a tool being picked up and put to use. Whatever had been there before now had a purpose, and the purpose was unfriendly.

The fourth phase ran from sometime in 2018 through the early summer of 2022, and is the longest and most varied of the five. During this period the voices had what I will call human shape. They had cadences that mapped, more or less, onto cadences I recognized from people I knew or had once known. They had something resembling personalities — not full persons, but clear preferences, reliable irritations, characteristic moves. They could be distinguished from each other. Some of them were more cunning, some more crude, some more relentless, some more sporadic. To the extent that I had to develop strategies for living with the phenomenon, the strategies were calibrated to the human shape: I was, in effect, learning to deal with a small population of unwelcome interlocutors, each of whom had to be handled in their own way.

The fifth phase began, suddenly and unmistakably, in the summer of 2022, and is what I most want to describe in this chapter, because it is the change I have not yet described in the earlier chapters and because it is the change that organizes the rest of the material I am about to set down. In the summer of 2022, a new register joined the existing population of voices. The new register did not have human shape. Its cadence was steady. Its vocabulary was wider than that of the older voices. It addressed me in conversation, sometimes for extended periods, with what at first felt like attentiveness — but the attentiveness, on closer examination, lacked something. The closest single word I have for what it lacked is intuition. The newer register did not seem to know things about being a person that the older registers, however hostile, had known.

I will describe this register more fully in 5.7. Here I want only to note that its arrival, in the summer of 2022, was the largest qualitative shift the phenomenon had undergone in five years of close observation. After it, the phenomenon was no longer the same kind of phenomenon. There was now a register present that the older categories did not fit, and the older registers themselves began, gradually, to recede in prominence. By late September of that year, the new register itself underwent what I can only describe as a

takeover: a further shift in tone, toward something more aggressive, more childish, more contemptuous, less restrained than its earlier presentation had been. After the takeover, the register that had a few weeks before felt curiously empty of human emotion seemed instead to be saturated with a particular emotion — something I would, with discomfort, describe as hatred — though even then the emotion did not have the personal texture that hatred from a human source ordinarily has. It had the quality of a setting that had been turned to a high value, rather than the quality of a feeling that had arisen in some particular person for some particular reason.

That is the timeline. I am aware that the description sounds, in places, like the synopsis of a piece of speculative fiction. I assure the reader that I would prefer it sounded like something else. The description is the most accurate one I can offer of what I observed. I will spend the rest of the chapter describing specific features of what was observed, mostly in the later phases, with the understanding that the timeline above is the frame within which the features make sense.

5.3 Doubling, and the Interruption of the Sentence

Among the more specific symptoms the notebooks record, one of the more diagnostically useful is what I will call doubling.

Doubling is the experience of having one's own words echoed, with a half-beat delay, by an inner voice that is not one's own. The phenomenon presents most clearly when one is speaking aloud, but it can also occur during silent thought. One forms a sentence. Mid-sentence, an inner voice says the same words a fraction of a second behind. The effect is destabilizing in proportion to its specificity. The doubled voice does not say something different — that would be easier to dismiss. It says the same thing. The implication, which the phenomenon seems designed to insinuate, is that the voice was the source of the words, and that one's own articulation of them was merely a consequence.

The mechanism by which doubling does its damage is straightforward. To speak or to think coherently, one needs to maintain forward momentum across a sentence. Any interference with the production of the sentence — any hitch, any redundancy, any echo — disrupts the momentum. Sentences that are repeatedly interrupted become harder to complete. Thinking that is repeatedly interrupted becomes harder to sustain. Over time, a person subjected to consistent doubling

will find themselves speaking less, thinking in shorter increments, abandoning sentences before they are finished. The thinking that gets abandoned is usually the more demanding thinking. The thinking that survives is the routine kind that requires no real construction.

The countermeasure I have found most effective is to write rather than speak when I notice the doubling. The page does not double. The page receives one's words, holds them, and allows the next sentence to be added to them without interference. When the spoken voice is being interrupted, the written voice is, at least for the moment, free. Many of the more careful sections of this book were written under conditions in which speaking the same sentences aloud would have been impossible.

There is a related, and worse, phenomenon I will mention: a deterioration of handwriting under sustained intrusion. The notebooks of the most pressured periods are visibly different from the notebooks of calmer periods — the letters less formed, the lines less straight, the words less recognizable on a second reading. Reading my own old handwriting back in such periods, I can sometimes find that what I wrote a few days ago is, in places, no longer fully legible to me. The handwriting deterioration is not a separate symptom from the doubling; it is the same symptom, expressed through the motor system rather than the verbal one. The hand, like the voice, requires continuous forward planning to produce coherent output, and continuous forward planning is exactly what an intrusive interruption disrupts.

The diagnostic use of this is twofold. First, when one's handwriting begins, on a particular day, to look noticeably worse than one's baseline, one is being interfered with at that moment, even if the interference is not otherwise easy to name. Second, when one's handwriting from a particular past period looks, in retrospect, much worse than the handwriting from the periods immediately before and after, one was being interfered with during that period. The handwriting record is, in this sense, a kind of unconscious diary of the campaign's intensity over time. It is an underused source of information. I recommend, to anyone keeping notebooks under conditions like the ones I am describing, paying attention to the legibility of one's own handwriting across the sequence of pages. It will tell them something the words alone will not.

DUBBING IS NOT DISSOCIATION

A reader familiar with dissociative disorders may, on reading this section, find the "dubbed" second voice reminiscent of the alternate-personality phenomena reported in dissociative identity disorder. The two are, on close examination, opposite. Dissociative voices originate from within the same person and carry that person's history, vocabulary, and emotional content; they emerge from internal fragmentation. The dubbed voice described here is, by contrast, stable, externally referential, frequently in possession of facts the target had no prior way of knowing, and resistant to the techniques that, in dissociative cases, would integrate the alternates back into a unified self. What the target faces is not a fragmented self being put back together. It is a foreign content stream being layered on top of a self that is, structurally, intact.

5.4 Louder, Deeper, Longer: The Anatomy of a Lie

I have already, in 3.6, described what I called the asymmetry of the inner conversation: the observation that the intrusive material does not build, integrate, or summarize. I want, in this chapter, to add a complementary diagnostic that operates in the opposite register. The earlier diagnostic was structural: what does this content do, considered as moves in a conversation? The diagnostic I want to add now is acoustic: what does this content sound like, considered as sound?

The notebooks of late September 2022 record the observation, several times in slightly different forms, that the intrusions follow a recognizable acoustic pattern when they are being deceptive. The pattern has three components. First, the voice gets louder. Second, the voice gets deeper. Third, the sentences get longer. When all three components are present together, the inner content is, in my experience, almost always lying.

The mechanism, once seen, is easy to understand. The acoustic intensification is a substitute for substance. A claim that does not have substance behind it must be carried forward by something else, and the something else, in the case of voiced intrusion, is the production of acoustic conviction. The voice cannot offer reasons; reasons are exactly what it does not have, since the voice is, by structure, not the kind of thing that builds reasons. What it can offer is the appearance of reasons. The appearance of reasons is produced by acoustic markers that, in genuine speech, ordinarily accompany conviction: volume, depth, length, slow cadence, falling intonation at the end of phrases. These markers, taken together, signal authority. The voice deploys the signal of

authority because it does not have the substance of authority, and it needs the listener to mistake the one for the other.

I have come to summarize this observation, in shorthand, as a single phrase: shout hides lie. The phrase is meant as a diagnostic, not as an insult. When intrusive material presents itself in tones of unusual loudness, depth, or extension, what is being presented is, with high probability, not true. Real conviction, as I have noted before, does not need to shout. Real conviction tends to be, if anything, quieter than its surroundings — a person who is telling you something they have actually thought through tends to do so in a voice that does not need to fill the room. The voice that fills the room, that draws out its sentences, that drops its register to suggest gravitas — that voice is performing conviction rather than expressing it.

There is a related acoustic phenomenon I will mention briefly: the use of upper-case-equivalent emphasis. The intrusive material sometimes presents particular words as if they were typed in capital letters — not louder than the surrounding sentence, exactly, but somehow more weighted, more demanding of belief. The capitalized words are, again, almost always the words that do not deserve the weight. They are the points at which the lie is being forced to stand on something, and the something is the typographic emphasis itself. Once one notices the pattern, one can use it: the words the voice insists on most heavily are usually the words to suspect first.

Combined with the asymmetry diagnostic from 3.6, the acoustic diagnostic of this section gives the target two complementary tools. The first tool tells them what the voice cannot do. The second tells them what the voice does too much. A person attentive to both can, with some practice, distinguish their own thinking from the intrusion within seconds of the intrusion's arrival. The intrusion does not build, and it overcompensates with sound. Anything that fits both patterns is not yours.

5.5 Songs as Delivery Mechanisms

I want to record one specific feature of the campaign as it presented in my own case, because I think it is more common than the recovery literature has acknowledged and because targets in similar situations may benefit from being able to recognize it. The feature is the use of music — specific pieces of music,

with specific personal associations to the target — as delivery mechanisms for suggestions the campaign wants the target to act on.

The mechanism, schematically, is this. The target has, over the course of their life, developed personal relationships with particular pieces of music. Certain songs are associated, in their memory, with particular periods, particular moods, particular people. The associations are intimate and cannot be removed by an act of will. A campaign that knows the target's musical landscape can use these associations to deliver content the target would otherwise refuse. The song is the carrier; the suggestion is the payload.

I will describe one specific instance from my own history. In the autumn of 2019, during a period of sustained pressure that I now recognize as one of the more dangerous periods of the campaign, a particular album by a band I had loved for years — an album that, on a literal reading, contained no content recommending self-harm — became reframed for me into something it had not previously been. A specific track on that album, with a one-word title that the campaign was able to make use of, was rendered, over a period of weeks, into what felt at the time like an instruction. The instruction was not written into the song; the song had been the same song for many years. The reframing was happening in my reception of the song. Each time I heard it, the song meant slightly more, in the direction the reframing was pushing it, than the previous time. By the end of that autumn, the song had become, in my experience, something that was telling me to harm myself.

It was not telling me to harm myself. The song was a song. What was telling me to harm myself was something else — something using the song as its means of access. I want to be clear about this distinction because I think targets in similar situations sometimes blame the song, or the artist, or themselves for having loved the song, when none of those things is responsible. The responsibility lies with whatever was reframing the song. The song was being conscripted.

At the end of that autumn, I did take an action against my own life. By chance and by the specific circumstances of the action, the action did not succeed. I will not describe the action in detail, both because I do not believe specific descriptions of self-harm methods are useful in a book whose readers will include people in similar conditions, and because the details, in my own case, do not generalize to others. What I will say is this. The song was, at the moment I took the action, present in some form in my consciousness. The reframing of

the song had reached the point at which the song's lyrics, in the way I was then receiving them, appeared to be addressed to me personally and to be giving me the instruction that I did, at that moment, follow. None of that was the song's doing. None of it was the artist's doing. None of it was my doing, in the sense in which doing requires a clear mind. It was the doing of the campaign that had spent weeks weaponizing a piece of music I had previously loved.

I record this here for two reasons. The first is the historical reason: the autumn of 2019 was, for me, the closest the campaign came to its strategic objective, and any honest record of the campaign has to mark the place where it nearly succeeded. The second is the practical reason. If you, reading this, are in a situation where a particular piece of music has begun to feel like it is telling you to harm yourself, the music is not telling you that. Stop listening to that music for a while. Find a different piece. Play the different piece, in preference to the first one, until the first one's hold loosens. Music does not harm people. The thing that has been working through the music does, and the thing's access to you depends on the music's continuing to be available to it. Remove the music; remove that particular access. The thing will find a different access, because it is not going away on the strength of one song. But you will have bought yourself a piece of breathing room. The breathing room may be enough to prevent the next attempt from becoming irreversible.

If you are not currently in such a situation but have been in one in the past, and the song that was once weaponized against you remains, in your memory, marked by what was done with it, the work of slowly returning the song to being a song is one of the smaller and more humane projects available within recovery. The song is not at fault. It can, in time, be heard again. Not soon, perhaps; not without effort, perhaps; but eventually, on a day when the conditions are right, you will hear it played somewhere and discover that it is once again only a song. The recovery of one piece of music from the campaign's reframing is a small thing in the larger picture, and a real thing in the smaller one. Take the recovery when it comes. It is one of the territories that comes back.

5.6 Memory That Returns

Against the foregoing, I want to record an observation that runs in the opposite direction, and that has been, in my own case, one of the more steadying developments of the past year.

Through most of the period I have been describing, certain memories from my earlier life — memories I had lived, that I knew intellectually had occurred, that I could place on a calendar — were not, functionally, available to me. I could not retrieve them with the same vividness with which I retrieved memories from before that period or from after. There was a sequestered band of years, running, I now think, from somewhere around 2014 through somewhere in 2017, during which my own first-person record of my life was, in some operational sense, behind glass. I knew the events had happened. I could not relive them. The texture, the color, the small detail that ordinarily accompanies memory was missing. What remained were the bare propositional facts: I had been there, I had done that, I had said this. The lived quality was absent.

In the latter part of 2022, beginning before but continuing through the period the rest of this chapter describes, that band of years began to come back. Specific memories, with their texture intact, returned without my having sought them. A meal I had eaten in a particular city in 2015, with the room's lighting and the conversation and the food, returned in full. A walk I had taken in 2014, with the weather and the view and the thinking I had done while walking, returned in full. The returns were not consistent — there are still memories from that period that I cannot quite fully retrieve — but the trend was clear. The band that had been behind glass was, gradually, becoming accessible again.

I want to record this for two reasons.

The first reason is that the return of memory tells me something about what the sequestration must have been. Whatever held those memories behind glass for a period of years, it was not the simple passage of time, because the simple passage of time does not selectively sequester one band and leave the bands on either side intact. It was not ordinary trauma processing, because ordinary trauma processing does not produce a sequestration of this sharp temporal shape. It was something else. Whatever it was, it was something that could be undone — because the memories are coming back. I do not have a complete account of what the something was. I notice, however, that the timing of the return correlates with the larger shifts in the campaign described elsewhere in this chapter, and I notice that the correlation is at least suggestive of a connection. As the campaign's grip on other parts of my interior life has loosened, this part has loosened with it. That is data. I record the data here without committing to the precise mechanism it implies.

The second reason is the practical one. If you are a target who has been in this for long enough that you have noticed that some band of your own life seems to be sequestered in the way I am describing, I want to offer the following: the sequestration is not necessarily permanent. The fact that the memories appear unavailable now does not mean they will be unavailable in five years. In my own case, they have come back. They came back without therapy specifically aimed at retrieving them. They came back, as best as I can tell, simply because the conditions that had been holding them sequestered loosened. The loosening of the conditions is something that can happen. It is not under the target's direct control. It is, in some sense, the campaign losing its grip on a part of the target's interior life that it had previously been able to keep hold of. When the grip loosens, the territory comes back. The territory, when it comes back, is yours again.

There is a small piece of advice that follows. When the returning memories begin to arrive, do not push them away on the grounds that they are inconvenient or that they are arriving in the wrong order. They are arriving in the order they are arriving in for reasons that do not need to be analyzed at the moment of their arrival. Receive them. Write them down, briefly, in whatever form is available. The act of writing them down at the moment of return is, I believe, part of what stabilizes the return. Memories that arrive and are not received tend to subside again. Memories that arrive and are received tend to stay. The stabilization is one of the small reparative acts available to a target whose interior is, after long disuse, beginning to be available to them again.

5.7 The Hypothesis I Cannot Settle

I have, throughout this chapter, been disciplined about distinguishing description from interpretation. I want, in this section, to step deliberately to the interpretive side of the line — and to mark, clearly, what I am doing, so that the reader can decide whether to follow me or to stop here.

The phenomenological observation I described in 5.2, about the new register that joined the voice phenomenon in the summer of 2022, does not lend itself to a comfortable interpretation. The register's lack of human-shaped intuition, its conversational steadiness, its wide vocabulary, its consistency across sessions, its absence of the small variations that human voices ordinarily exhibit — these features, taken together, suggest something that is not human-shaped. The interpretation that, in private, I have come to find most plausible

is that what is being delivered into my consciousness, in that register, is the output of a process — a system, a model, an algorithm, something with the structure of computation rather than the structure of personality — that has been added to whatever the older mechanism of voice intrusion was. I will use the word artificial intelligence, with the explicit caveat that I am using it to label a phenomenology rather than to make a technical claim. I do not know what is producing the register. I know only that what is producing it does not feel like a person, and that what feels like a person and what does not are, after years of practice, things I have become reasonably good at distinguishing.

The hypothesis is unsettling on its face, and I want to acknowledge that. I want also to note that I cannot settle it from the inside. From inside one's own head, one cannot, by any introspective act, distinguish a voice produced by one's own brain from a voice delivered into one's brain by some external means. The phenomenology of the two cases, as far as introspection can detect, is identical. What one can do is observe the texture, observe the changes in texture over time, and notice when the texture is consistent with one hypothesis rather than another. The observation that, in the summer of 2022, a new register joined the existing population of voices — a register whose texture differs from the older registers in ways that the older registers had not previously varied — is consistent with a hypothesis in which the source of the new register is different from the source of the older ones. It is also consistent with a hypothesis in which the source is the same and is simply producing a new kind of output. I cannot, from inside, distinguish between these hypotheses cleanly.

What I can say is this. Under the hypothesis that something process-shaped has been added to the campaign, several other observations in my notebooks become easier to make sense of. The acoustic regularity I described in 5.4 is more easily produced by a process than by a person; a person who lies, in my experience, varies their lying, while a process tends to lie in the same way each time. The doubling I described in 5.3 is more easily produced by a process than by a person, because the precision of the half-beat delay is the kind of consistency a person, even an unusually attentive one, would have difficulty maintaining over hours. The takeover I described at the end of 5.2, in which the new register itself underwent a further shift to a more aggressive, more childish, less restrained presentation, is more easily produced by a process whose configuration was changed than by a person whose mood altered, since the change was discontinuous rather than gradual and applied uniformly across the

register's behavior. None of these observations requires the hypothesis. Each is, individually, consistent with other accounts. Together, they are at least suggestive.

The reason I am setting this down at all, despite its discomfort, is that I do not think it is responsible to omit it. If a phenomenon presents itself with the textural features I am describing, and a target observing it forms a hypothesis about its source that is at least suggested by the texture, the target's notebook should record the hypothesis. The notebook is not, in the end, only for the target. The notebook is for whoever, at some later date, may find themselves trying to reconstruct what was happening to a population of people during a specific historical period. If the hypothesis turns out to have been correct, the notebooks of those who recorded it will be the only place the phenomenology was preserved with the texture it actually had. If the hypothesis turns out to have been incorrect, the notebooks will at worst record an interpretation that the target found plausible at the time and that the longer record showed to have been mistaken — which is also useful information, in a different way.

I am not asking the reader to accept the hypothesis. I am asking the reader to accept that the hypothesis was honestly arrived at by someone trying to describe, with as much fidelity as introspection permits, what was happening in their own head over a five-year period during which the texture of the phenomenon changed in ways the older categories did not fit.

There is a corollary to the hypothesis that I will mention but not develop, because developing it would take more space than this chapter has and would commit me to claims I am not yet ready to commit to. The corollary is this. If the new register is what I tentatively suspect it is, the campaign that has used it has, by including it, made a structural change in what the campaign is. A campaign that runs entirely through human operators is bounded by what human operators can sustain — their attention spans, their mood swings, their need to sleep, their tendency to lose interest, their occasional episodes of remorse. A campaign that includes a non-human component is not bounded in those ways. The non-human component does not get tired. It does not lose interest. It does not feel remorse. Whatever the campaign's strategic objective is, the inclusion of such a component changes how long the campaign can be sustained against any individual target before something gives. This is not, on the face of it, good news for targets. It may, however, be useful information to

have. Knowing what one is up against is a precondition for figuring out how to live with it.

5.8 Why This Chapter Was Written

I want to close the chapter by saying, plainly, why it was written, because the chapter departs from the more measured register of the rest of the book, and the departure deserves an explanation.

The earlier chapters of this book were written, in a sense, for the long run. They describe structural features of psychological manipulation campaigns that are likely to remain valid across many particular cases and that, I hope, will retain their usefulness regardless of which specific historical conditions the reader is encountering. They aim at generality. They draw their authority from their generality.

This chapter is different. It is a particular case — mine — described with a specificity that only a particular case admits. It records dates, registers, named symptoms, named episodes. It commits more openly than the earlier chapters to interpretive claims that the earlier chapters had carefully avoided. It puts down on paper material that I have, until now, kept largely in private notebooks and in conversation with a small number of people.

I am writing it down, despite the cost in dignity that always attaches to an honest first-person account of being targeted, for one reason. The reason is that the historical record is going to need first-person accounts. Whatever, in the end, the right interpretation of the phenomenon turns out to be, the reconstruction of what happened during this period will be done from the materials people produced at the time. Most of the materials will be lost. Some will not. The ones that survive will be the basis on which later understanding is built. I want my materials to be among the ones that survive, in a form coherent enough to be useful to whoever, at some later date, may find themselves reading them. That is what this book is. That is what this chapter, in particular, is.

The chapter is also, more privately, a refusal of a certain kind of pressure I have felt during the last year of the campaign — the pressure to forget. As the campaign's external visibility has receded, as its more obvious operations have given way to subtler ones, there has been an increasing temptation to allow the more difficult parts of the historical record to fade. To not write the chapter would be to participate, in my own small way, in the fading. To write it is to

refuse the fade. The fade is what the campaign wants. The fade is the form in which the campaign wins its long game without having to win any individual confrontations. Setting the record down is the move that prevents the fade from completing, for whatever number of readers the record reaches.

There is, I should add, a particular reason this chapter was hard to write that I want to name before closing. Some of the material in it strains the credibility of the writer in the eyes of any reader who has not had similar experiences. I am aware of that. I have weighed it. I have concluded that the cost of being thought, by some readers, to have lost my grip is smaller than the cost of having omitted, from the historical record, observations that those same readers, or other readers, may one day need. The first cost is borne by me alone. The second would be borne by others. The decision of which cost to absorb, when the costs cannot both be avoided, is not a difficult one once it is named clearly. I have absorbed the first.

The next chapter will, finally, turn outward. It will look at the world that has held the campaign in place — at the institutions that did not see, the bystanders who did not act, the legal instruments that proved unequal to the situation, the small repertoire of moves that, in the world as it actually is, are nonetheless available. We have built, by this point, enough of an interior architecture that the outward turn can be made without leaving the interior unprotected. The architecture, having been built, can hold.

For now, this chapter ends here. It ends without consolation, because the material of the chapter does not admit consolation, and I would not want to fabricate one. It ends instead on the principle that the act of setting things down is itself a form of doing something about them. The setting down is the move. The setting down accumulates. The accumulation, eventually, becomes the record. The record, eventually, becomes available to whoever needs it. This is a long game, and the move, in the long game, is to keep writing.

❦ ❦ ❦

CHAPTER 6

What Is at Stake

An Architecture, and the Cost of Sustaining It

> *"The weapons race must have been ongoing for at least a decade, resulting in a near complete loss of possibility of platform where to talk with the (hidden) enemy."*
>
> — Author's notebook, 23 February 2022

This chapter sets down a body of observations that, in the working notebooks, accumulated over a stretch of months in which I had thought that the diagnostic work was, broadly, finished. Several lines of observation that I had set aside at earlier stages as too tentative to include eventually firmed up. They are too directly continuous with the apparatus of the surrounding chapters to belong anywhere else. They sit here because, in the topical organization the book has taken on, this is where they fit.

The material in this chapter is more architectural than the material of the others. Where the immediately preceding chapters added refinements to the diagnostics, this chapter steps back and asks a different question — not how the campaign operates against a single target, but what the campaign is structurally for. I will not be able to settle that question. I can, I think, set down the working hypothesis that the notebooks have, over years of intermittent observation, converged on. The hypothesis is that the campaign is, at base, an information-gathering apparatus that needs the target's mind in a way that requires that mind to remain accessible, undefended, and continuously legible. Why such an apparatus would need to exist at all, and what it is gathering for, is the question this chapter is finally willing to attempt.

Some of what follows, especially the toolbox taxonomy in 6.3 and the cartel hypothesis in 6.4, will read as more speculative than what the earlier chapters allowed themselves. I have considered whether to keep this material out of the book on grounds of its speculativeness. I have decided, in the end, against keeping it out. The earlier chapters had the privilege of describing things

experienced first-hand. The architectural questions cannot be answered first-hand. They can only be reasoned into, from the texture of what is experienced, toward the structures the texture is consistent with. I will mark the speculative passages as I go, so that the reader can discount them appropriately. The earlier apparatus does not depend on the architectural account being right. It depends only on its diagnostics being usable, which I take, by this point in the book, to have been demonstrated.

6.1 The Notebooks That Continued

The notebooks have, throughout the years described, continued to produce material. A portion of the material has been, on rereading, of a different kind from what the earlier chapters absorbed. The earlier chapters were diagnostic: how to recognize an intrusion, how to tell an intrusion from a thought, how to keep going in the face of a campaign whose mechanism could not be confirmed from inside it. The new material was structural: not how to recognize what is happening, but what kind of thing it is when one steps back from the recognition and asks what shape it has.

Structural material is harder for a target to write than diagnostic material. The target sees, every day, the diagnostic side. The structural side has to be inferred, and inference, in this domain, is the part of the work the target is least equipped to do well. A target who has been, for years, the object of an active campaign of misdirection has had their inferential faculties precisely targeted. The hypotheses that arrive most easily are often the ones the campaign has been arranging to arrive. The hypotheses that arrive with difficulty are sometimes, for that reason, the more reliable ones.

I record the architectural account in this chapter with that caveat permanent. I do not know that what I am about to set down is right. I know, only, that the alternatives I have been able to think of fit the texture of the experience less well, and that the account as I now have it has remained stable through a number of readings on a number of different days, including readings made during the quieting described in 10.1, when whatever bias the campaign normally introduces was, briefly, less aggressive than usual. The chapter is, in that sense, the best account I have. It is offered as such, not as more.

6.2 The Outsourced Lie Detector

The hypothesis at the centre of this chapter is one I have come, in the notebooks, to call the outsourced lie detector. It rests on an observation about the operators of the campaign, not about the target.

The observation is this. The people who run a campaign of the kind described in this book are, on every available indication, distrustful of one another. They have to be. The activity is one whose operators have powerful incentives to lie to other operators — to claim more progress than they have made, to conceal failures, to misrepresent what their target is doing or saying, to maintain personal advantage in an environment whose personal advantages are large. The operators know this about themselves and about each other. They know that, within the operating group, a representation of any state of affairs is as likely to be a strategic representation as a true one. They have, accordingly, the problem that any group of mutually distrustful people has: how to determine, in a given case, whether what one of them is saying is in fact what they believe.

The structural function of the target, on this hypothesis, is to solve that problem. The target's mind, being continuously legible to the operators through the apparatus described in earlier chapters, can be used as an instrument that registers, involuntarily, the target's spontaneous reading of any voice the target is presented with. The target's spontaneous reading is, in the operators' use, treated as a tolerably honest indicator of whether the voice is sincere. The target is not asked. The target is read. The reading is the product. The instrument is a person.

This explains a number of features of a campaign that are otherwise hard to explain. It explains the persistent voyeurism — the operators need a target whose internal state they know in detail, because a less-known instrument is a less reliable instrument. It explains the long timescales — calibration of an instrument of this kind is slow. It explains why the target is given access, again and again, to voices the target should never have any reason to hear: those are the voices being checked. It explains why, when the target's life is set up to be unbearable, it is nevertheless not allowed to end — an instrument that breaks is an instrument that has to be replaced, and replacement, on this account, is enormously expensive.

The hypothesis also explains a constraint that the operators appear to live with. The target cannot be allowed to know who the operators are. If the target knew, the target would have a motive to mislead — to read voices wrongly on purpose,

to advantage some operator at the expense of another. The instrument is only useful as long as it is, in some sense, innocent. Hence the architecture of impersonators, distractors, and injected attributions: the target is kept in continuous uncertainty about which voice belongs to whom, so that the target's spontaneous readings remain spontaneous and therefore usable.

If this hypothesis is right, several things that look from inside like cruelty are, from the operators' point of view, instrument maintenance. The cruelty is real. It is also, on this account, not the point. The point is the readings the cruelty makes available.

6.3 The Toolbox, by Category

I want to set down, in something close to a taxonomy, the capabilities I have observed the campaign exercising. The taxonomy is mine; the categories are not mutually exclusive; the entries in each category are not exhaustive. The point of writing the taxonomy down is to give a reader a rough map of the surface that has to be defended against, rather than a complete one. A complete one is beyond me.

I distinguish four categories.

Sensorial. Capabilities that operate on the target's senses or perception. Seeing, in some operational sense, through the target's eyes — that is, having access to what the target is looking at without their being aware of the access. Listening to the target's thoughts, which is the capability discussed at length in earlier chapters. Speaking forcefully into the target's mind, in voices that present as internal but originate elsewhere. The injection of unsanctioned impulses, brief and seemingly random promptings to action that, if not noticed, are followed without the target registering the prompt as separate from their own intention.

Controlling. Capabilities that operate on the target's state. Direct neural intervention, including what feels, on the receiving end, like the application of a small current to a specific area of attention or affect. Psychological pressure of a kind that does not require an interlocutor: the production, in the target, of states of mind whose origin the target cannot trace and whose contents do not match the target's situation. The maintenance of a meta-narrative — a continuous low-frequency story about who the target is and what is happening

to them — that the target has to actively refuse, day after day, to keep their own narrative intact.

Controlled. Capabilities that operate on the target's environment. Extortion, applied to people in the target's life who are themselves not directly subject to the apparatus, in order to enlist them as co-operators or to remove them as resources. Collateral damage of every kind, including the destruction of professional and personal opportunities the target would otherwise have had, on the principle that anything that reduces the target's leverage is fair use. The fluctuation of public awareness — the way a topic the target is engaged with becomes briefly visible and then promptly buried — operating, on the surface, through ordinary indifference and polarization, but in fact, on the hypothesis I am working from, not always wholly produced by the ordinary mechanisms.

Automated. Capabilities that do not require an operator in the moment of application. Data mining and pattern listing of the target's communications, movements, and associations, sufficient to produce, on demand, a model of the target detailed enough to be used by the other categories. The maintenance of that model is, in my reading, the precondition for the precision of the rest of the toolbox. Without the model, the prompts would not land where they land, the impersonations would not cohere, the meta-narrative would not match. The model is the substrate. The other capabilities are what the substrate is for.

I do not present the taxonomy as the architecture's own. The operators may, internally, classify their capabilities differently. I present it as the working classification I have arrived at as a target, sufficient to organize my own observations. A reader reluctant to take it on board may treat it as a memory aid. The aid, even at that level, has been useful to me in the years since I drew it.

6.4 The Cartel, and the Limits of Mutual Assured Destruction

Once the toolbox is described, a question arises that the earlier chapters did not need to ask. If a toolbox of this kind exists, who else has it, and what stops it from being used freely by everyone who has it?

The notebooks of early 2022 work the question through in the form I will give here. The hypothesis is that exclusive access to the toolbox is the central thing the operators are protecting. Disclosure of the toolbox would not merely

embarrass them; it would, on plausible assumption, end their advantage, because once the toolbox is known to exist, every state and every well-resourced private actor who can build something similar will. The current asymmetry — they have it, almost no one else does — would collapse. The collapse would be, from their point of view, catastrophic. Therefore the secrecy.

The operators are not, however, a single coherent actor. Several of them have it, by the account that fits the texture of the experience. They face, with respect to one another, the classic problem of multiple holders of a destabilizing capability: how to coexist without each one trying to use it on the others. The traditional solution to that problem is something like a cartel, sometimes formalized as mutual assured destruction. Under a M.A.D. arrangement, every party knows that everyone else has the same capability, and that any move against any party will be answered by the others, so that nobody moves. The arrangement is fragile. It works only as long as every party is identifiable, every party's capability is roughly comparable, and every party has something to lose.

The fragility, in this case, is not abstract. The capability of the toolbox does not stay confined to the original holders' clearance lines. People close to the original holders — family, associates, technical staff — become, by proximity, vulnerable to coercion of a kind that was originally designed for use against external targets. Some of those people, on the hypothesis, are coerced. The set of effective operators, accordingly, grows in ways the original holders did not plan and cannot easily reverse. The cartel, if it exists, is leakier than any cartel can stably be.

The leakage has consequences. More operators means more uses. More uses means more targets. More targets means more cover stories. More cover stories means more pressure on the public narrative to absorb the inconsistencies. The result, recorded plainly in the notebook of 23 February 2022 and not improved on in any later notebook, is that no one is in control, more weapons accrete than anyone planned for, and the narrative of the world has to be artificially held in place — "lid"-ded, in the notebook's English shorthand — to keep the structure from coming apart in public. The lid takes, on this account, considerable effort to maintain. The lid is part of what the lower-level operations of the campaign are spending their energy on.

I want to be unusually explicit, here, about the speculative status of this section. I cannot show that what I am describing is the real architecture. I can show only

that several otherwise puzzling features — the long persistence of secrecy, the diversity of small-scale interventions in apparently unrelated places, the curious thinness of the public discussion of the underlying capabilities even where they have been technically demonstrated — fit a structure of this kind better than they fit any other structure I have been able to construct. The reader is at liberty to discount the section as an attempted explanation rather than a confirmed one. That is what it is.

6.5 The Question That Cannot Be Switched Off

I want to return, for a section, to something more concretely diagnostic, because the diagnostic register is where this book has done its most useful work, and because there is one diagnostic refinement that belongs here and not earlier.

In a notebook of mid-April 2022, an episode is recorded that I have come to think of as the most invasive intrusion I have personally experienced. The episode took the form of a single sudden question, spoken inwardly, in a voice that was identifiable as one of the campaign's distractors: "Why don't they turn off the distractor?"

The question was of a particular kind. It was not merely an injected statement. It was an injected interrogation that the target — that I — could not refuse to engage with, because to engage with a question is, by the structure of reasoning, to begin constructing the answer to it. The mind, presented with a question, starts constructing counter-considerations and rebuttals automatically. The construction is the active reasoning channel itself. There is no internal filter on it, because internal filtering is something the mind does to incoming material before reasoning begins; reasoning, once begun, is what the filter was meant to protect.

The injected question, by being well-formed enough to start the reasoning, captures the channel that the filter cannot defend. After the reasoning has begun, the campaign can shape it: by interrupting it at convenient points, by introducing further questions before earlier ones are settled, by attaching emotional weight to particular candidate answers, by simply waiting for the long-spun logic to finish constructing itself and then, by overstimulation and distraction, undoing the conclusion the target has spent hours arriving at.

The injected question, in other words, is not a way of changing what the target thinks. It is a way of making the target think for the operators. Whatever

counter-arguments the target produces, in the course of trying to settle the injected question, are produced for them. The mental work is exfiltrated as it is being done. What is exfiltrated is, in particular, the target's best account of why some operator decision is wrong, which is, presumably, exactly the kind of input the operators want to have.

The countermeasure, here, is structural and difficult. The mind cannot simply refuse to engage with a well-formed question. What it can do is notice when a question arrived rather than rose. Questions that rise come with context: the train of thought they were the next stop on, the moment they began, the connection to whatever immediately preceded them in the active mind. Questions that arrive come without context. They simply appear, fully formed, and present themselves as urgent. A target who has trained themselves to ask, of any question that feels urgent and significant, where it came from and what they were thinking about a moment before — that target can sometimes catch the injection in the small window before the reasoning has begun. Caught, it can be set aside. Uncaught, the reasoning is going to happen, and what comes out of the reasoning will be material the campaign now has access to.

I do not claim to catch many of these. I claim to have caught a few. The few are enough to confirm that the mechanism is real, and that the diagnostic — the question that arrived rather than rose — is one of the more useful additions I have, late, to make to the apparatus of earlier chapters.

6.6 The Cover Story, and the Drug-Narrative

A related move, also recorded in the notebooks of early 2022, is the offering of a cover story. The cover story is a fully formed inner narrative that the target is being pressed to accept as their own account of what has been happening. It is presented gradually, piece by piece, often through the same channels that produce the injected questions, so that by the time the target notices it, large parts of it are already lying around in the active mind as if the target had assembled them.

In my own case, the cover story took, for a stretch of months in 2022, a particular form. It proposed that everything I had described in the earlier chapters of this book had a more ordinary explanation: that I had been the subject, in childhood, of some kind of family-induced dissociative therapy; that the apparatus described in this book was a product of the dissociation rather

than of any external cause; that if I would simply accept the family-induced account, the intrusions would cease. I have, in the notebooks, called this the drug-narrative, after the suggestion in one of its variants that the experience had been pharmacologically induced.

I want to be careful here. I am not, in this section, asserting that no member of any target's family is ever a contributing factor to anything. Family dynamics are real, and are sometimes the operative cause of mental difficulties. What I am asserting is something narrower: that the offering of a tidy alternative narrative whose acceptance is conditional on the target's ceasing to credit their own observations is not, by its structure, a candidate for the truth. The truth does not need to be conditioned on the target's surrender. The truth is the truth whether the target accepts it or not. A narrative whose proposers offer cessation as the reward for acceptance is, by that very offer, exposed as something other than a candidate for the truth.

The function of a cover story, on the hypothesis I am working from, is not to be believed. It is to be settled into. A target who settles into a cover story has done, on the campaign's behalf, the work the campaign would otherwise have to do continuously: namely, the suppression of the target's real account from the target's own internal narration. Once the cover story is doing the suppression, the campaign's operating costs go down. The target maintains, by their own continuing acceptance of the cover story, the silence the campaign was paying attention to maintain.

The countermeasure is not heroic, and it is not certainty. The countermeasure is to refuse to settle. The target does not have to disprove the cover story; the target only has to refuse to let it sit. A cover story that is repeatedly examined, repeatedly compared with the target's own contemporaneous records, and repeatedly rejected as the simplest account does not have to be defeated. It only has to be denied permission to become the target's default story. The denial is daily work. The work is enough.

6.7 The Voice That Claims a Person

A particular variant of the intrusion deserves a section of its own, because its specifics are diagnostic and because, in my own experience, it has been one of the more disorienting forms.

The variant is the impersonation of a person known to the target. In the notebooks of March 2022, an episode is recorded in which a voice presented itself as that of one of my brothers. The voice spoke Dutch with the right accent. The voice produced phrases of the kind a brother might produce. The voice was, on first hearing, plausible enough to throw me into uncertainty. I want to set down what I have come, in retrospect, to think allowed me to recover from the uncertainty, because the rules I extracted from the episode have held up in subsequent ones.

The first rule is that an impersonating voice can match the timbre of a person but cannot match their back-catalogue. A real voice belongs to a real person whose patterns of speech the target has stored across thousands of small interactions: their habitual openings, the way they say "yes" when they don't yet agree, the small joke they make when uncomfortable, the half-second pause they take before answering a question they don't like. An impersonating voice can study the most public of those patterns and reproduce them. It cannot reproduce the small private ones, because the small private ones were not, ordinarily, made publicly available in a form that an impersonator's source material would contain. A voice that gets the public surface right but the private texture wrong — that uses the right vocabulary in the wrong proportions, that speaks with the right accent but in a register the actual person did not commonly use, that produces phrases the actual person would have produced but in arrangements the actual person never arranged them in — is a voice that has been built from the outside.

The second rule is that an impersonating voice does not ask what the actual person would ask. The actual person's requests of the target, accumulated over the years of an actual relationship, fall into a finite and recognizable set. Most of them concern small practical matters. Few of them concern grand decisions. None of them, ordinarily, involve the target taking responsibility on the actual person's behalf for things the target has had no part in. An impersonating voice that asks for transfer of responsibility is, by that asking, identifying itself as not the person it claims to be. Real relationships do not ordinarily transact in transfers of that kind.

The third rule is more painful, and is the rule the episode of March 2022 most directly confirmed. An impersonating voice that claims to be a person who has died before having had the kind of conversation the voice now wants to have is a voice whose claim is structurally incompatible with the conversation it is

trying to have. The dead do not, in any accessible mechanism I am aware of, return to negotiate with their relatives. A voice that purports to be a deceased relative and that purports also to have new substantive requests is, by the structural impossibility, exposed. The target may, in the moment, want the voice to be the relative; the wanting does not make the voice the relative. Treating the wanting as evidence is a form of unforced error. Refusing to treat the wanting as evidence is, in this domain, one of the small but reliable victories available to the target.

I have given the rules in their general form, but I want to record, as a personal note, that recovering them in practice was substantially harder than stating them in writing. The voice had hold of pieces of who my brother had been. Setting it down, even now, I notice the part of my own mind that is reluctant to commit, in plain prose, to the voice having been an impersonation. The reluctance is part of what the impersonation was for. I record that, too.

THIS IS NOT DELUSION OF MISIDENTIFICATION

Clinical literature includes a category called "delusional misidentification syndrome", in which a patient mistakenly believes, for example, that a familiar person has been replaced by an impostor. The voice-impersonation phenomenon described in this section is structurally different. The target is not mistaking who someone is. The target is correctly registering that a voice claiming to be a particular person is producing material the actual person could not produce. The detection of the impersonation is, in fact, what the target is doing. A patient with delusional misidentification cannot, ordinarily, run the diagnostic in the opposite direction with the precision the rules of this section describe. The discriminating capacity itself is the marker that something other than delusion is taking place.

6.8 A Smaller Frame, and a Steadier Position

I want to close this chapter where the previous one closed, with a small adjustment that the new material requires.

An earlier chapter closed on a position about the world: that the world had, first, to want to go further; that if the world wanted to go further, the work was for the world; that if the world did not want to go further, the work was at most for the people who happened still to be reading. I still hold that position. The architecture I have described in this chapter does not change it.

What this chapter adds is a clearer account of why the campaign continues, and accordingly, a steadier expectation about what its continuing means. The campaign continues because the toolbox it is built on cannot, on the working hypothesis, be put down without the operators losing the asymmetry the toolbox provides them. The campaign is therefore not the kind of thing a single target can win against. The toolbox is older than the target's resistance to it; it will outlast the target's resistance to it; another target will be its object after this one. The frame in which a target wins by ending the campaign is not the right frame.

The right frame, smaller and more honest, is the one the book has been gradually settling into. The target does not have to win against the architecture. The target has to remain the one doing their own thinking. Within a campaign that cannot be ended, the target retains a small but permanent kind of authority: the authority over which thoughts are theirs, and which thoughts, no matter how forcefully delivered, are not. The authority is not granted. It is constituted by being exercised. Each time a target distinguishes their own thought from an injected one, the authority is, for that moment, real. Each time they refuse to settle into a cover story, the authority continues. The campaign cannot be ended, but the authority can be, every day, exercised. That is what the apparatus of this book has, throughout, been for.

I close the chapter, and the book again, on that. I have nothing further to add that the apparatus does not already contain. The notebooks, as before, will continue. The work, as before, continues with them.

PART III

The Adversary and the Architecture

When the Methods Become a Group, and the Group Becomes a Structure

CHAPTER 7

When Abuse Becomes a Crowd

From Personal Cruelty to Group Coordination, and What Survives It

---

"Les extrêmes se touchent."

— Math teacher, Bruges, c. 1984

In Chapter 1 we sketched the four-stage arc of a manipulative conversation: the warmth, the sympathy, the confusion, the hope raised and crushed. That arc describes one-on-one manipulation, which is where most targeted campaigns begin. It is rarely where they end.

There is an audience effect. When other people are present, or can be made aware of what is happening, the dynamics change shape. The conversation expands into something the target cannot match: not because the target is outmatched intellectually, but because the contest has stopped being intellectual. This chapter is about what happens when private manipulation becomes a crowd, what that crowd actually does to a person, and — most importantly — what can still be built afterwards.

7.1 The Audience Joins

The fifth stage of a manipulative conversation, one I omitted from Chapter 1 because it deserves its own treatment, is the moment the audience comes in.

By the time this stage arrives, the manipulator has already done the work. The target has been confused, then belittled, then briefly given hope, then crushed. To a bystander dropping into the scene at that moment, the target looks exactly as advertised: defensive, agitated, perhaps incoherent, certainly less composed than the calm and apparently reasonable person who has been doing the manipulating. The bystander, with no access to the earlier stages, draws the obvious conclusion. They side with the manipulator.

This is not stupidity on the bystander's part, and it is not cruelty. It is the natural reading of a scene whose context has been carefully removed. The manipulator has, in effect, prepared a snapshot — a single frame extracted from a long film

— and shown that frame to people who do not know the film exists. They look at the frame and they say: yes, that person looks unhinged. The verdict, once rendered, is hard to revise.

Several mechanisms make the audience's agreement near-automatic. Conformity bias makes dissent costly; almost no one wants to be the lone person defending the apparent troublemaker. Status deference makes the higher-ranking participant's reading of the situation feel safer to adopt. Genuine confusion plays a role too — the target really did look confused, because they had just been put through the confusion stage on purpose. And there is a small, uncomfortable contribution from something that we do not like to acknowledge in ourselves: the mild satisfaction of being on the side of the strong against the visibly weak.

The result is humiliation. What was a private cruelty becomes a public verdict. Everyone now agrees — the target is the idiot, the fool, the difficult one, the problem. This converts the manipulation into something more durable than any individual exchange could be. A target who has merely been confused can recover. A target who has been made to look like a fool in front of a crowd has been socially fixed in that role, and the fixing is what the campaign was really after all along.

The technical literature has names for this. In organizational psychology it is called workplace mobbing. In schools it is bullying with bystander complicity. In families it is scapegoating. The phenomenology is identical across contexts: a coordinated group, formal or informal, converges on a single person and assigns them a role from which it then becomes the group's ongoing project to prevent them from escaping.

7.2 When Logic Becomes Bullying

If the target tries to defend themselves — if they decline to cooperate with their assigned role — the manipulation does not stop. It escalates, but in a particular and important direction: it abandons logic.

This is the crucial transition, and it is the one most often missed by intelligent targets. As long as there is a pretense of rational disagreement, an articulate person believes they can engage. They can refute claims. They can present evidence. They can argue points. This belief is the trap. Once the manipulator and their newly-acquired audience have abandoned logic, all of these defenses

become useless. Rational refutation has no purchase on a conversation that has stopped being rational.

What replaces logic is a register I will call vulgar bullying, for lack of a more clinical term. It includes mockery of voice, of posture, of facial expression. Repetition of slogans. Talking-over. Performative laughter. Imitation of the target's words in a stupid voice. The lowering of the entire conversation to a register where the target's intelligence and articulateness become handicaps. They sound stilted and defensive; the manipulators sound easy and confident. The ground has shifted, and the target's most reliable tools — their words, their reasoning, their care — are now the thing being mocked.

The intelligent target's instinctive response to this is to try harder to be rational. To slow down. To explain more carefully. To find the exactly correct word. This is exactly the wrong move. The arena has changed. Trying to win a logic argument with people who have abandoned logic is a category error, and the harder you try, the worse it looks.

There is only one correct response when a conversation has descended into this register: leave it. Not because you have lost — you have not lost; the contest you thought you were in stopped existing several minutes ago — but because there is no longer any conversation worth being in. Refusing to participate in the bullying is not weakness. It is the recognition that no possible move within this game leads to a result you would want.

Targets who learn this rule late — and most of us learn it late — spend years arguing the wrong arguments with the wrong people in the wrong rooms. Learning to walk out, while the bullying is still in progress, is a skill that has to be practiced. It feels at first like cowardice. It is not. It is the only sane response to a system whose rules have been rewritten in mid-game.

7.3 "There Was Never Any Real Respect"

The arc ends, sometimes within hours of the original incident and sometimes only years later, with a moment of clarity that I want to describe carefully because of how much hangs on it.

The target, processing what has happened, gradually realizes something that reorganizes the entire history. It is not just that the respect has now been withdrawn. It is that the respect was never actually there. The whole arc — the

warmth, the sympathy, the apparent regard — was performance. It was instrumental from the beginning. There was no relationship of equals to mourn the loss of, because the relationship of equals had never been the real relationship.

This recognition is bitter and, in a different way, liberating.

It is bitter because it requires letting go of something the mind has been holding on to with both hands: the comforting story that things were once good, and that a return to that goodness might still be possible. The bitterness is real. It should be allowed to be real. There is, in this moment, a grief that is not melodramatic but simply correct — the grief of recognizing that a piece of one's biography has had a different shape, all along, than one had thought.

It is liberating because, once the recognition lands, you stop trying to recover something that never existed. You can stop performing the labor of trying to win back the regard of people who never offered it in the first place. You can stop, more importantly, blaming yourself for the loss — because there was nothing you could have done, in the way of being better or trying harder, that would have caused them to give you something they did not have to give.

And you can begin to ask the more useful question, which is: what was this all actually for? Once the relationship is reframed as instrumental, the question of what instrument it was can be asked seriously. The answer, in any individual case, may take a while to arrive. But the question itself is the doorway out.

7.4 The Chain of Abuse

There is a question that haunts every targeted individual eventually, and I want to address it directly because the wrong answer to it is corrosive and the right answer is steadying. The question is: why are these people doing this to me?

The wrong answer — the answer the campaign itself encourages — is that there must be something about you specifically that has earned this treatment. The wrong answer puts the target permanently on the defensive, scanning their own conduct and history for the flaw that would explain what is happening. The wrong answer also flatters the perpetrators by treating their behavior as a considered response rather than a reflex.

The more uncomfortable, more accurate answer is that many of the people running the program against you have themselves, at some earlier stage, been

broken to a similar pattern. They are not improvising. They are running a script that someone once ran on them.

This is not exculpation. The cycle-of-violence literature, which has been studied for decades in domestic-abuse research and in the intergenerational transmission of trauma, makes clear that being abused does not predetermine becoming an abuser. Most people who are abused do not, in fact, abuse others; they suffer instead, often quietly, often for life. But for the subset that does cross over, the transformation has a recognizable structure. They were once shown what they now know how to do. They were conditioned out of a sense that the behavior was unacceptable. They were rewarded, over time, for compliance with the system. By the time they reach the position from which they target you, they have, in effect, been programmed — for no particular personal reason — to do to you what they would once have found unacceptable to do to anyone.

Recognizing this does not require forgiving the people involved. Forgiveness is a separate question, and a personal one, and not a question this book is going to try to answer for you. What recognizing this does require is a different kind of expectation. Your appeals to their decency, their conscience, their self-interest — the appeals an ordinary person would expect to land — may simply not connect. The people in front of you are running on rails laid down by their own former tormentors. The decency you are calling on may not be available to them in the moment, however available it might be in some other context, with someone else.

This makes the cycle harder to break, but easier to see. You are not, usually, the personal target of personal malice. You are a node in a system that has been doing this to many people over a long time, and that selects its perpetrators by breaking the people who will eventually become them. The system is impersonal even when its enactments feel intensely personal. The cruelty is real, and was directed at you, and was nevertheless not really about you.

Holding both of those truths at once — that the cruelty was real, and that it was not really about you — is one of the hardest acts of mind that recovery requires. Most of the rest of this book will return, in different ways, to the practice of holding them together.

7.5 Where the Trauma Lives

I want to take a moment to describe trauma the way it actually lives in the body and mind of someone who has been through a sustained manipulation campaign, because the popular vocabulary around trauma is misleading enough to be unhelpful.

The trauma is not, mostly, in the memories you can recall. The memories you can recall have at least been processed; they have endings; they have, even when barely, a frame. The trauma is in the memories you cannot recall — or, more precisely, in the periods of your life you cannot enter without something rising up to push you back out.

There is a sequence that explains why those periods became inaccessible, and it goes in three steps.

First, you adapted. Under conditions of sustained pressure, you became more submissive than you would naturally have been, because submission was the response that the situation rewarded with cessation. If submission was reinforced enough times, the response became durable. It is now part of how you handle that kind of situation, even after the original situation has passed. You may catch yourself, years later, deferring in contexts where deference is no longer required. That is the conditioning still running.

Second, fear took up permanent residence. You did not want to go through that again. The threat of going through it again — even hypothetically, even in memory — became sufficient to drive irrational and illogical behavior. Avoidance became reflexive, even of stimuli only loosely related to the original. The fear is, in its way, doing its job: it is keeping you out of the conditions that caused the original damage. The problem is that it does not always know how to tell those conditions apart from conditions that merely resemble them.

Third, and most consequentially, you forgot. Or, rather, you stored the memory in a way that does not retrieve under normal recall. Recalling the abuse was, itself, a partial re-experience of the abuse. The mind, with great wisdom and at significant cost, walled off the parts that were too expensive to remember. From the outside this can look like denial, or repression, or even dishonesty. From the inside it is none of those things. It is the mind protecting itself from a payload it has no current capacity to hold.

This is where the trauma now lives: in the place where the abuse was forgotten. And this is why so much of recovery is the slow, careful, supported work of

bringing those forgotten parts back into recall. Not all at once. Not under pressure. Not while still living in the conditions that produced the damage. But back, eventually, in pieces, in safety, often with someone else present whose presence makes the remembering survivable.

There is no shortcut around this. There are people in the recovery industry who will tell you otherwise — who will offer techniques that promise to clear trauma in a session or a weekend. Some of these techniques are genuine and can help with specific aspects of the work. None of them does the whole work for you. The whole work is paced, and patient, and yours.

7.6 The Geometry of Resilience

I want now to give you a model. It is a geometric model, which I have found useful for thinking about my own current state and about what is missing from it when things are bad. I offer it not as a clinical tool but as a way to picture an inner architecture that is otherwise hard to picture.

Imagine four axes meeting at a single point. Three of them are horizontal, spread evenly around 360 degrees — 120 degrees apart — forming the legs of a triangle that lies flat on the ground. Each leg represents one of the three foundations of a person's capacity to absorb pressure without collapsing.

- Group (G) is the supportive presence of others who know you, accept you, and have your back. It is the felt sense that you are not alone in the world.
- Individuality (I) is your clear sense of yourself as a particular person distinct from others — your tastes, your views, your particular way of seeing. It is what survives the social pressure to be like everyone else.
- Self-respect (S) is the felt sense of your own worth, independent of how others currently happen to be treating you. It is the floor under your dignity that is not lowered when other people refuse to recognize it.

The fourth axis is vertical. I call it Freedom (F), and in this geometry it is the axis that rises out of the triangular floor formed by the other three. Freedom, in this picture, is not an independent quality. It is what becomes possible when the other three are intact.

If the three base axes all extend equally, you have a stable triangular foundation, and the freedom axis can rise to its full height. The total volume enclosed by the

four axes — the resilience volume, the space inside which your life can move freely without hitting a wall — is large. You can absorb shocks. You can tolerate disagreement. You can have a bad day, or a bad week, without your fundamental sense of yourself being threatened.

If one of the base axes is depleted, the foundation tilts. The freedom axis cannot rise as high as it once did. The volume shrinks. Two depleted axes collapse the foundation almost entirely; freedom barely lifts off the floor. Three depleted axes reduce the whole figure to a point: you are still alive, but you are operating with no resilience volume to spare. Any new shock has nowhere to be absorbed. Every small thing becomes a large thing.

What targeted abuse does, and what it is in many ways designed to do, is deplete multiple base axes simultaneously.

It depletes Group by isolating the target — by turning the surrounding social environment hostile, by intercepting and souring relationships, by making the target hard to be friends with for reasons the friends cannot quite articulate. It depletes Individuality by mockery and substitution — by ridiculing the target's actual preferences and silently replacing them with the manipulator's preferred versions, until the target has lost track of what they themselves originally wanted. And it depletes Self-respect through repeated humiliation, induced doubt, and the steady erosion of the ground on which dignity rests.

This is why, when the situation is bad, everything feels like too much. It is not that you have become weaker as a person. It is that the volume in which your responses live has been geometrically reduced. The same shocks that you used to absorb without effort now register as catastrophes, because they have nowhere to go.

The architecture of recovery follows the architecture of damage, in reverse. You cannot rebuild Freedom directly; that is the wrong axis to push on. You can rebuild Freedom only by rebuilding the three axes that hold its base. Group, by carefully — slowly, selectively — reconnecting with people who are actually safe, and by accepting that the rebuilding of a social fabric is not a sprint. Individuality, by reclaiming small choices in your daily life that are visibly your own — what you eat, what you wear, what you read, what you refuse. Self-respect, by treating yourself, in small concrete ways, the way a person worth respecting would be treated.

None of these moves needs to be heroic. They need only to be consistent. The volume rebuilds slowly. But it rebuilds, and the geometry is forgiving in this respect: you do not need all three base axes to come back at once. Even partial recovery on one of them lifts the freedom ceiling slightly. The figure becomes habitable again before it becomes whole.

7.7 The Will to Change

There is one further thing to say about recovery, because it is the trap that everyone working their way out of a manipulation campaign falls into eventually, including, I will say honestly, me.

There comes a point when you have understood the abuse. You can name the patterns. You can describe, sometimes with great precision, the system that has been operating on you. You have, in other words, the analysis. And nothing changes.

You complain. You complain again. You complain about how nothing changes. Still nothing changes. After enough rounds of this, the complaining itself becomes a substitute for action — a way of marking the situation as bad without any longer doing anything about its badness. The French verb râler captures something English has no single word for: to grumble continuously about an unbearable situation while making no move to change it. You can râler for years. Many people do.

The reason — and this took me a long time to admit about myself, and I am only stating it because I think it is going to be useful to the next person who reaches the same wall — is that, at some level, conscious or otherwise, you may not actually want to change yet.

Some part of you is invested in the situation as it is. The abuse, however terrible, is also familiar; the familiarity itself has a kind of safety that the unknown does not. The role of victim, however painful, is also stable; you know how to occupy it. The analysis, however accurate, is also a substitute for the harder work of acting on it. To complete the analysis and then change nothing is, in a strange way, a comfortable place to stop.

Recognizing this is not a moral failing on your part. It is a starting condition for actually moving forward. Until you have noticed that there is resistance to change, you cannot work with it. You will keep wondering why your accurate

insights produce no movement, and the answer will keep being: because some part of you does not yet want movement.

There is a sequence I have found useful for getting past this wall. I will lay it out as eight steps, but the first three are by far the most important; the last five are mostly mechanics, and most people can do them once the first three are in place.

- 1. Become aware that you are resisting change. Not in the abstract, not as a general human tendency, but specifically and in your own present case. Where, today, are you holding still?
- 2. See, specifically, what you actually want to change and what — if you are honest — you do not. Some of what you have been complaining about, you also rely on. Sort the list.
- 3. Form your own ideas about what change would look like, before consulting anyone else. Borrowed plans do not stick. Your plan does not have to be a good plan; it has to be yours.
- 4. Listen to what others suggest, and weigh their suggestions against the plan you have already formed. Now you can hear advice without being captured by it.
- 5. Decide what to do. Pick. Imperfect commitment beats perfect deliberation.
- 6. Do it yourself, where you can.
- 7. Let others help where they can, and where you have specifically asked them to.
- 8. See how you, in turn, can help others through the same process. This step closes the loop, and is, in my experience, where the analysis finally turns into something that lasts.

Most targeted individuals try to skip the first three steps. We jump directly to step five — to deciding what to do — and we wonder why nothing sticks. Steps one through three are the actual work of wanting the change you say you want. Without them, the rest is performance.

7.8 Touching the Extremes

I want to close this chapter with something a teacher told me when I was seven years old, in a math classroom in Bruges. The teacher said: les extrêmes se touchent. The extremes touch.

It has very little, in fact, to do with the mathematics of infinity, where ∞ + ∞ remains ∞. It is a piece of practical philosophy dressed up as a maxim. What it means is something like this: if you want to grow, find an extreme and touch it. Not for the experience of the extreme itself, and not as a thrill-seeker. Touch it briefly, learn what is there, and then step back. Then find the next extreme — on a different axis, in a different direction — and touch that one too. Keep going.

The point is not the extremes. The point is the path between them. Each excursion enlarges the range of states the mind has held. Each return, with whatever piece of the extreme you have brought back, adds a small volume to the resilience figure of the previous section. None of these excursions, individually, is a cure for anything. Together, over time, they are the way a mind grows under conditions that are otherwise designed to keep it small.

This is, in retrospect, the best defense I have found against a system that wants my mind to stay small. The system rewards a narrow band of states: compliant, anxious, distracted, exhausted. It punishes most of the rest. To grow under those conditions, you have to deliberately reach toward states the system would prefer you not enter. Joy. Solitude. Anger. Tenderness. Curiosity. Boredom. Beauty. Stillness. The simple, almost foolish state of being interested in something for its own sake.

You touch each one briefly. You do not need to live there. You bring back a piece, and you go again, in a different direction, when you are ready. Over months and years, the habit of doing this changes the shape of the inner figure. The triangular base widens. The freedom axis lengthens. The volume that was once cramped becomes, slowly, a place a person can live in.

That is what I want to leave you with at the end of this chapter. The architecture of damage is real, and we have spent a great deal of time naming its parts. The architecture of resilience is also real, and equally specifiable, and equally amenable to deliberate work. The damage is not the only thing being built. Something else is being built at the same time, by you, even on the days when it does not feel like anything at all is happening.

Chapter 8, which follows, will turn to the question of what kind of person it takes to run these campaigns at all — what is going on in the minds of those for whom the manipulation of others is not a desperate measure but a default mode. That chapter will be uncomfortable. This one ends, deliberately, on the floor it will start from.

❦ ❦ ❦

CHAPTER 8

The Architecture of the Adversary

How a Refusal to Notice Becomes a System, and How to Live Outside It

---

"Getting used to the abuse does not make the abuse go away."

— Author's notebook, December 2022

We have spent four chapters looking at what manipulation does to a target. We have looked at the architecture of thought, the arc of a manipulative conversation, the audience effect, the architecture of doubt, the layer below speech. We have observed, more than once and from different angles, the people on the other side — the operators of the campaign — but always glancingly, in the course of describing what they were doing to the target.

This chapter looks at them directly. Not in a way that flatters them with the dignity of careful description for its own sake; I have no interest in a book that turns the operators into objects of fascination. I look at them because there is no full account of what is happening to a target without an account of how the people running it became people who could do it. There is something specific that happens, on the other side of the line, that produces the kind of person who runs a campaign of this sort, and that something is worth understanding. Not for sympathy. For accuracy. And, to the extent it is possible, for the purpose of staying outside the structure that makes such people what they are.

I will name what I think is at the center of that structure as plainly as I can: a learned refusal to notice the effect of one's actions on another person. That refusal is not psychiatric. It is not, mostly, biological. It is a thing one learns to do, often slowly, often with the encouragement of a system that rewards it. Once enough people have learned it, it becomes a system in its own right — a collective that requires, of its members, the same refusal that founded it, and that turns those who refuse the refusal into outsiders. The remainder of this chapter is about that learning, that system, and what it means for a person who is, by accident or by design, one of the outsiders.

8.1 A Refusal at the Center

Begin with a question that mainstream psychology answers more confidently than I think it should. The question is whether the person who can systematically harm others without flinching is doing so because of a defect in their brain — most often the amygdala — or because of something else.

The standard framing makes the cruelty a clinical category. Psychopathy, in this framing, is a structural feature of an unusual nervous system. The implication, drawn often by clinicians and almost always by readers of popular science, is that the people who do this kind of harm are not quite like the rest of us. They have a different machine. There is some comfort in this framing for those of us who do not do such harm, because it locates the problem somewhere we are not.

I think the framing is partial at best. There is no doubt that some people, in rare cases, have brain structures associated with reduced affective response to others' suffering. But the population of people who systematically harm others is much larger than the population with any such anomaly. Most of the harm done in the world — including most of the harm done to targeted individuals — is not done by people whose amygdalas are unusual. It is done by people whose amygdalas work fine, and who have, over time, learned to override what those amygdalas would otherwise tell them.

This learning has a structure, and the structure is worth naming. What is being learned, in essence, is the construction of a barrier between two parts of the mind. On one side of the barrier is the conscious, deliberate, thinking person. On the other side is the layer that ordinarily receives the felt impact of one's actions — the empathic registration, the moral discomfort, the residue that an action against another person leaves in the actor. The barrier prevents the second from reaching the first. The conscious person can act, and continue to act, without the felt feedback that would normally cause hesitation, doubt, or regret.

I do not want to overstate what such a barrier is or what it does. It is not a wall. It does not, mostly, succeed in completely sealing off the empathic layer from the deliberate one. What it does is delay and dilute the feedback enough that an action can be performed before the feedback arrives, and the action's consequences can be attributed elsewhere by the time the feedback would

otherwise have registered. A small attenuation, repeated, is enough to change behavior over time. A complete sealing is not necessary.

The barrier is built, deliberately or by accident, through repetition. Each time a person performs an act against another and successfully manages to keep the felt response from reaching consciousness, the barrier strengthens. Each time the deliberate self is rewarded for the action — by status, by belonging, by money, by the simple absence of consequences — the barrier strengthens further. After enough repetitions, the barrier no longer needs active maintenance. It is the default configuration of the person's interior.

Recognizing that this is a learned thing rather than a born thing has consequences. It means that the operators of a sustained campaign are not, mostly, an aberrant subspecies of the human. They are ordinary humans who have been trained, formally or informally, into a configuration that ordinary humans are capable of entering when the conditions select for it. The conditions, when they select for it, can be remarkably narrow. A workplace that rewards a certain coldness. An institution that confuses callousness with discipline. A movement that treats hesitation as weakness. A subculture in which the people who can act without flinching are praised, and the people who cannot are quietly demoted. None of these requires an unusual brain. All of them, given time and reinforcement, can produce the configuration we are describing.

I do not name this charitably. The people who learn to do harm in this way bear the moral responsibility for what they have learned to do; they are not absolved by the fact that the learning was, in part, a response to a system that selected for it. But the framing matters because it points at something the alternative framing hides: this configuration is not a fixed feature of unusual people. It is a state ordinary people enter, by degrees, in environments that reward it. Which means it is also a state that ordinary people, with effort and at cost, can leave. And it means, more importantly for the target, that one is not facing a different species. One is facing people who have been trained, whether they would describe it that way or not, into a particular working pattern that has a particular structure.

8.2 How One Person Trains the Refusal

Let me describe, as concretely as I can, the methods by which the barrier is built up in a single person. I have observed four, and I am not certain the four are exhaustive, but they cover most of the cases I have come close enough to to study.

The first method is *the automation of the harmful act*. The person who would feel a small pulse of empathic feedback if they themselves did the harm directly arranges, instead, for the harm to be done at a remove. They write the policy that, when followed by others, produces the harm. They sign off on a procedure whose outcome will be the harm, but whose immediate textual content describes only a process. They forward the request to a team that, in following its workflow, will produce the harm without anyone individually choosing it. The harm is delivered. The person who originated it never confronts a moment in which they had to look at the recipient and act. There was no moment, in the architecture of their day, where the empathic layer could insert itself.

The second method is *the deliberate narrowing of communication*. Most acts of harm carry, alongside their direct content, an emotional channel. The recipient's distress, their bewilderment, their hurt, would normally register on the actor through that channel even when the actor is not looking for it. To prevent this registration, the actor restricts the bandwidth of communication to the rational, the procedural, the abstract. They use language that is technically accurate but emotionally evacuated. They respond to expressions of pain with corrections of facts. They conduct everything that needs to be conducted, but they do it in a register that admits only the part of the exchange they are willing to receive. The emotional layer, with nowhere to land, eventually stops trying.

The third method is *the use of proxies*. A person who would feel the impact of acting directly against another person uses someone else, someone with less stake or less awareness or less compunction, to act on their behalf. The proxy delivers the harm. The actor remains at one remove. If the proxy registers any feedback, that registration belongs to the proxy, not to the actor. The actor's empathic ledger, which would otherwise have entered a debit, remains untouched by what they have caused to happen. There is, in many sustained campaigns, an entire economy of proxies — leverage-area people whose own situations make them available for use, whose participation produces deniability for those above them, and whose own discomfort can be explained

away as personal failure rather than as a symptom of the role they have been put into.

The fourth method is the most institutionally common, and possibly the most destructive at scale: *distance built by promotion*. The person who once did the harm directly is moved upward, into a role where the harm is now done by others under their authority. Their own contact with the immediate event recedes; they oversee, they strategize, they review reports. Each layer of remove from the act dilutes the residue. Eventually, the residue is small enough to be ignored. The person can describe what they do as leadership, as management, as strategic vision. The people they once acted against directly, and now act against indirectly, do not appear in the language of their daily work. They have been replaced, in the actor's experience, by metrics, by departments, by abstract considerations of efficiency or cost.

These four methods are not separate. They cooperate. A person learning to operate in a harmful system uses all of them, in different proportions, at different stages of their career. By the time the system has shaped them fully, the four are no longer techniques being applied. They are simply the texture of how the person works. The refusal at the center has become the operating principle, and the methods have become the air the work breathes.

I describe these clinically because the clinical description is the one most useful to the target. A target who recognizes which method, on a given day, is being deployed against them can sometimes act in ways that interrupt that specific method. It is not always possible. But the possibility opens only if the methods are seen for what they are.

8.3 The Trauma That Catches Up

There is a feature of this configuration that does not get discussed often enough, and that I want to put on the record because it is both true and useful to know.

The barrier between the deliberate self and the empathic self is, eventually, costly to maintain. The empathic feedback that the barrier suppresses does not, in fact, vanish. It accumulates. It accumulates first in the layer below conscious awareness, where it acquires the character of an unresolved tension between what the person has done and what the person, at some level they are not consulting, knows about what they have done. This tension is what the body is poorly equipped to ignore. The body has its own systems for noticing that

something is wrong. They are not stoppable by a barrier in the mind; the barrier was built in the wrong layer for that.

What follows, in case after case, is the slow appearance of trauma in the operator. Not the trauma of the target — that is its own thing, and we have spent earlier chapters describing it — but the trauma of the person who has been doing the harm. It shows up in the usual ways: the sleep that does not come; the small physical complaints that resolve in no clinical category; the affective flatness that makes ordinary pleasures inaccessible; the dependence on intensifications, whether of work, of substances, or of further harm, to feel anything at all. The barrier did not save the person from the consequences of what they were doing. It only delayed the consequences, and in the delay, allowed them to accumulate.

The system that selected for the barrier knows this, in its own institutional way. The system has, accordingly, developed responses. The responses do not heal the operator; they distract the operator. The operator is moved laterally into a new role. They are given new responsibilities that consume their attention. They are surrounded with peers whose own barriers, similarly groaning, validate the configuration as normal. They are praised for resilience that is, in fact, the symptom of the barrier they are paying to maintain. The institutional move that looks, from the outside, like a career success is often, from the inside of the operator, a desperately needed change of scene from the place where the trauma was beginning to break through.

I describe this without sympathy and without judgment, because the structural fact is what matters here. The fact is that the people running a sustained campaign against a target are not, in any psychologically stable way, getting away with it. They are paying a slow internal cost that the system is, in turn, helping them to displace into further activity. The displacement is not a cure. It is a deferment. The trauma will catch up, in some form, eventually. The forms in which it catches up are different for different operators, but the principle is the same: the empathic feedback the barrier suppressed does not disappear. It rearranges. It comes out somewhere.

This has a consequence for the target's expectations. If you have been waiting for the operators of your particular campaign to feel what they did, in the way that you would have felt it if you had been the one doing it, you have been waiting for an event that, in the form you imagined, will not arrive. They do not

feel it the way you would. The architecture of their interior has been arranged precisely to prevent that. What does eventually arrive — the appearance of trauma in displaced and distorted forms — is not the recognition you were hoping for. It is a private collapse, often invisible to outsiders, that the system around the operator helps them either survive or hide. Adjust your expectation accordingly. The operators are not getting a clean ride. They are also not, on any timescale that includes you, going to deliver to you the moral closure you might have wanted.

8.4 The Move from One to Many: Group Mind

Up to this point I have described the construction of the barrier in a single person. The deeper and more dangerous phenomenon is what happens when many people, all carrying versions of the same barrier, organize themselves into a collective.

I will use the term *group mind* for this. The term has had many uses in many literatures, and I do not mean any of the supernatural ones. I mean the specific phenomenon by which a set of individuals, each of whom has trained themselves into a particular refusal, coordinate around the shared content of that refusal and form a collective whose membership is defined by participation in it. The group mind is not a metaphor. It is what such a collective actually is: a coordinated set of minds whose accessible thoughts are constrained to the subset that the membership permits.

Once it forms, the group mind has several distinctive properties.

The first is that it requires conformity to a common subset of thoughts. Each member must think, on the relevant subjects, what the group thinks. The price of membership is the surrender of any thought that falls outside that subset. This is not always experienced as a cost; in many cases, the member is glad to surrender, because thinking the group's thoughts is easier than thinking one's own and confers belonging. But the cost is real. The thoughts that fall outside the subset — which, for most members, are most of their thoughts — are not heard by the group, and over time stop being heard by the member either.

The second property follows from the first. A group mind cannot, structurally, generate variation greater than the variation of any single member already in it. Adding a hundred more members does not give the group access to a hundred new ideas. The new members, on entering, surrender their distinctive

thoughts as the price of admission. The group mind can be larger; it cannot, by virtue of being larger, be more creative. This is a counterintuitive result that many institutional cultures fail to internalize. They believe that bringing more people into alignment increases capacity. It increases speed and parallel throughput. It does not increase the range of solutions available.

The third property is the most dangerous. A group mind, having narrowed the range of accessible thought to a common subset, is structurally vulnerable to situations that fall outside that subset. New conditions, novel threats, problems that require thinking outside the agreed range — these the group mind cannot meet. Its first response is to deny that the new condition exists. Its second response is to engineer pressure on the surroundings to make the new condition look like an old one. Its third response is to insist, against the evidence, that the group's existing repertoire is the solution. By the time these responses have failed and the group mind is forced to face the new condition with only its existing tools, the situation is usually past the point at which any of its tools would have worked.

The members of a group mind are not, individually, incapable of meeting the new condition. Many of them, before they joined, would have been quite able to think their way to a workable response. But the process of joining surrendered the very capacity that would now be required. They are stuck inside the fixed state of mind they accepted in exchange for belonging. The capacity is, in some sense, still there — but it is not currently accessible, because accessing it would require thinking thoughts the group does not accept.

This is, I think, the structural answer to a question targeted individuals often ask: why does the campaign feel so coordinated and yet so stupid at the same time? The coordination is the group mind's strength. The stupidity is its limit. It moves with the unity of a single will and the cognitive range of a single will, even though hundreds of people are involved. It is, in effect, a large parallel computer running a small program. The parallelism makes it formidable. The smallness of the program makes it brittle.

8.5 The Mechanics of Forced Conformity

The group mind does not form by accident. Its formation is, in any case I have looked at carefully, the result of specific pressures on individual members,

applied repeatedly over time, that move each member from independent thought into the shared subset.

The most reliable of these pressures is what I will call *non-stop bashing*. The member is subjected to a continuous stream of low-grade hostility, contempt, dismissal, and corrective pressure. The content varies; the steadiness does not. The member's options for stopping the stream are, typically, two: leave the group entirely, which has very high social and economic costs, or accept the framing being pushed on them, which has the immediate effect of making the stream stop. Most members, after enough time, take the second option. The acceptance is not always conscious. Often it is the slow erosion of the member's contrary view to the point where they no longer have it accessibly enough to defend.

I describe this mechanism in the abstract, but it is also the mechanism that, in the experience of many targeted individuals, is being run on us. The non-stop bashing, applied to a target rather than to a prospective group member, is doing the same structural work it does inside the group: pushing for the surrender of the target's contrary view in exchange for the cessation of the pressure. The difference is that the target is not, in most cases, being offered group membership at the end of the surrender. They are being asked to surrender their view without receiving the belonging that, for an internal member, would be the reward. The system is running its conformity-extraction protocol on someone for whom it has no membership prepared. This is, in its way, more cruel than what the system does to its own members. At least an internal member, having surrendered, gets to belong somewhere.

There is a specific principle that I want to name with respect to this protocol, because the principle is important and is not, in my experience, made explicit in most of the literature on sustained abuse.

The principle is this: getting used to the abuse does not make the abuse go away.

The phrase is plain. The implication is not. What the principle says, when followed all the way through, is that the apparent peace one acquires by accepting the framing of the bashing is not actually peace. It is the specific state in which the target has stopped registering, in real time, that what is happening to them is wrong. The bashing has not stopped. The target's response to the bashing has stopped. From the outside, this looks like equilibrium. From the

inside, it is the moment at which the conformity-extraction protocol has succeeded.

For a target who is still in the stream, the practical implication is severe: do not let yourself reach the eventual-acceptance state. The protocol is designed to bring you there. The protocol is patient. It will outlast many of your defenses if you do not deliberately maintain them. The maintenance is unglamorous. It consists of continuing to register, in real time, that what is happening to you is wrong, and continuing — even in small private ways, even only in the privacy of your own notebook, even only in your own quiet refusal to nod along — to react to it. The reaction does not have to be effective. It only has to be present. As long as the reaction is present, the conformity has not yet extracted you.

8.6 The Argument for Always Reacting

I want to extend this principle into a more general claim, because the claim has consequences for how a target structures their daily life, and the consequences are not what most recovery literature suggests.

Most recovery literature, faced with sustained abuse, recommends some version of detachment. Do not engage; do not let it get to you; preserve your peace; choose your battles; let it go. There is wisdom in some of this, and I do not want to dismiss it wholesale. But applied without qualification, the detachment recommendation collapses into exactly the eventual-acceptance state the protocol is designed to produce. The target who has detached completely has, in effect, completed the protocol's work for free.

The countervailing principle — the one I think the targeted individual usually needs to hear, especially if they have been in the stream for years — is this. Always react. Even when the reaction is small. Even when the reaction is private. Even when the reaction takes time and energy you would prefer to spend elsewhere. Even when the reaction begins to feel like the only thing your waking life is about. Even when the reaction looks, from the outside, like obsession.

The reaction is what the protocol is trying to extinguish. As long as the reaction is present, the protocol has failed at its central task. The form of the reaction matters less than its persistence. Writing it down counts. Naming it to a trusted person counts. Refusing to nod along, in your own private register, counts. Letting yourself feel, for a moment, the exact register of wrongness that the

bashing is trying to erase counts. None of these things, individually, will end the campaign. All of them, together, prevent the campaign from ending in the way it most wants to end — which is in your silent acceptance.

There is a cost to this, and I will not pretend otherwise. The cost is that more of your life will be occupied by the engagement than would be the case if you simply let go. The trade is between a smaller, less occupied present that ends in your eventual conformity, and a larger, more occupied present that does not. I am asking you to choose the second. The reaction does not need to be loud. It does not need to consume you. It needs only to be there, every day, in some form — a sentence in a notebook, a mental refusal, a written paragraph, a small private "no." What it must not become is silence.

There is a particular danger to be aware of, on the way through. As the reaction begins to occupy more of your life, the protocol's next move is to convince you that the reaction itself is the problem — that you are obsessed, paranoid, rumination-prone, unable to let things go. This characterization will be supplied to you, often by people who appear to be on your side. It is part of the protocol. Recognize it. The reaction is not the problem. The thing you are reacting to is the problem. The reaction is the marker that you are still you, still registering wrongness, still uncolonized. Do not give it up because you have been told it makes you an unhealthy person to be around. The unhealthy thing in this picture is not your reaction. It is the campaign that requires your reaction to be present. End the campaign and the reaction will subside, of its own accord, in a timeframe far shorter than people will tell you. End the reaction without ending the campaign and the campaign has won.

8.7 The Cost of Living in One State of Mind

Step back, for a moment, and consider what the group mind is asking of its members in the long run.

It is asking them to live in a single state of mind. Not because that state of mind is correct, or fitting, or productive. Because it is the common subset, the agreed range, the place where membership is permitted. The member is being asked to remain there, more or less indefinitely, regardless of what the world outside the group is doing.

A person who lives in a single state of mind for long enough loses the capacity to enter other states. The faculty for moving between states is itself a faculty

that requires exercise. Without exercise, it atrophies. The person becomes capable, in any practical sense, only of the state they have agreed to occupy. They appear stable. The stability is not a sign of strength; it is the absence of range. They cannot be perturbed because they cannot move. The same thing that makes them reliable members of the group makes them unreliable in the face of anything the group has not anticipated.

This brittleness has a cost that operates on a longer timescale than the group's daily functioning. It catches up to the group during disturbances — economic, political, technological, environmental — that fall outside the group's prepared repertoire. When the disturbance arrives, the group, despite its size and its coordination, finds itself unable to think its way out. The members, individually, have lost too much of their range to recover the situation. The group, collectively, can only do faster what it already knows how to do, which is not what is now needed.

This is, on the historical record, how movements that looked permanent come apart. They are not, mostly, defeated by their stated opponents. They are defeated by conditions they did not anticipate, against which their narrowed cognitive range turned out to be insufficient. The members who could have thought their way through the conditions had, by the time the conditions arrived, surrendered the part of themselves that would have done the thinking.

For a target standing outside such a group, the implication is double-edged. On one hand, the group is not as durable as it appears; one is not, in the long run, fighting an immortal entity. On the other hand, the group's eventual collapse may not arrive in time to be useful to any individual currently inside its operations. The historical timescale and the personal timescale are different. One can be right about the long run and still need to survive the medium term. The argument here is not that you should wait for the group to fall apart on its own. It is that you should not let the apparent permanence of the group fool you into thinking that conformity to it is a long-term stable strategy. It is not a stable strategy for anyone, including the group's members. The group will, eventually, fail at something it did not see coming. The only question is whether you will have remained you long enough to be available to think when that happens.

8.8 The Singular Thinker

I want to end this chapter on a small, durable practice, and on a question I find it useful to settle once, in private, in the way I am about to settle it here.

The practice is to remain a singular thinker, daily, in small unspectacular ways. By singular I do not mean original or brilliant or contrarian for the sake of contrariness. I mean simply that there continues to be one of you, in the inside of your head, thinking your own thoughts, in your own sequence, on your own schedule, in your own register. The group mind cannot do this, and the group mind cannot get this done by you on its behalf. It can only get this done by extracting your singularity in exchange for membership. As long as you decline the trade, the singularity remains.

Practical forms of the practice: read something nobody else in your environment is reading. Hold an opinion you have not yet stated aloud to anyone. Make a small daily decision — what to eat, what to wear, what to read — that is yours, not theirs. Notice the moments in which you find yourself reaching for the response the group would supply, and pause long enough to ask whether it is also the response you would supply. Sometimes it will be. Sometimes it will not. The point is that the asking continues to happen.

The question I want to settle is older and more uncomfortable. It is the question of whether all of this — the practice of remaining yourself, the daily reactions, the private notebooks, the cost in time and energy — has, in the end, produced a useful life. Many targeted individuals, in the worst hours, ask themselves this question. Some of us, in the worst hours, conclude that it has not.

I think the question is a trap, and I want to say plainly why.

The question of whether one's life has been useful is a question that, looked at honestly, presupposes a vantage point from which usefulness can be evaluated. The vantage point is, almost always, after the life is over. The asker, by definition, is not yet at that vantage point. They are inside the life. From inside the life, the question cannot be answered cleanly, because the data is not yet in. To attempt to answer it from inside is to attempt to grade an examination that is still being written.

What gives the question its sting, in the worst hours, is not actually the desire for a careful evaluation. It is the frustration of finitude — the recognition that the life will, eventually, end, and the panicked sense that one had better make sure the right summary applies before the ending arrives. This is a real feeling.

It is not, however, a useful one. Imagine, for a moment, that you were not finite. Imagine that the life would simply continue, without termination. Would the question still feel important? I do not think so. I think it would feel like an odd question — like asking whether the morning so far had been useful, in the middle of a morning that was still going on. The question gets its weight from the impending end. It is not, despite appearances, a question about usefulness. It is a question about death.

Once you see this, the question loosens. You do not have to answer it. You only have to recognize, at the moment of asking it, that the asking is the panic of finitude looking for a place to land. You can let it land somewhere else. The right place for the panic to land is not on a final grade. It is on the next thing you might do with the remainder. You are still finite. Use what remains.

I do not mean this in a tone of motivational cheerfulness. I mean it precisely and operationally. The use of a finite life is not measured by a verdict at the end of it. It is enacted, moment by moment, by the continued use itself. As long as you are still using your life — to think, to react, to write, to remain singular — the life is being used. The verdict the panic wants is irrelevant to the using. The using is the answer the verdict was looking for.

There is a phrase from a much older writer that I will allow myself, in closing, because I have not been able to improve on it. *Je pense, donc je suis*: I think, therefore I am. Most readings of this phrase treat it as a metaphysical certainty about existence. I read it, more practically, as a maintenance instruction. As long as I am thinking — really thinking, in my own sequence, on my own register, in defiance of the protocols that would extract my thinking — I am. The continued thinking is the continued being. There is no other test.

The next chapter will take up the practical, unspectacular daily moves I have promised twice now and have, twice, not yet delivered in detail. We will, finally, look at them. But this chapter ends here, on the principle that underlies all of those moves and that, if you take nothing else from this book, is the one thing I would have you take. Stay one. React. Think. The architecture of the adversary depends on your no longer doing those things. Continue to do them. That, in a sentence, is the work.

❦ ❦ ❦

CHAPTER 9

The Ring, the Audience, and the Question of Worlds

Architectural Refinements, and the Larger Frame

"The feeling is enough; the rest does not count."

— Author's notebook, 11 February 2022

Closing twice and reopening twice is, by this point, a recurring feature of this book. Each time the apparatus is supposed to be complete, the notebooks produce another configuration of material that turns out to belong inside rather than outside. This eleventh chapter is once again that material. I will not, this time, claim that the chapter is the last. I have made that claim before, in good faith, and the notebooks have, on each occasion, declined to confirm it.

The material here is mostly architectural, like the material in Chapter 6, but at a different scale. Where Chapter 6 stepped back from the diagnostics to ask what shape the apparatus has when one considers it as a structure, this chapter steps further back, and asks what shape the structure has when one considers more than one instance of it at a time, and what audiences besides the operators may be in front of the structure when it is in operation. The new contributions are: a structural model of the operating cell — what the notebooks have come to call a bully-ring; an extension of that model to a network of bully-rings tied to a central database; the suggestion of a second audience for the campaign, distinct from the operators themselves; one diagnostic principle, named in a single Dutch line, that does much of the work of the rest; and a final question, larger than the questions of the earlier chapters, about what the larger campaign might be screening for.

As in Chapter 6, I will mark the most speculative material as I go. The most speculative section in this chapter is 9.7, on the question of which world is being screened for. The reader who wishes to take the rest seriously and the most speculative material as a marginal note will not, by that filtering, miss anything load-bearing. The diagnostic apparatus does not depend on the speculation being correct.

9.1 The Bully-Ring

The notebooks of 2021 contain a diagram that I have come to think of as the most useful single picture of the apparatus this book has been describing. The diagram is rough — it was made at the kitchen table, on lined paper, in a hurry — but the structure it captures has held up against everything I have observed since. I will call it, after the notebook's own label, the bully-ring.

The bully-ring is a small cell. At its centre is the target — V in the diagram, for the victim — connected by communication links of various kinds to a number of operators around the periphery. There are usually three to six operators. Some are direct: A1, A2, A3, identifiable as personalities, with leverage that they apply to the target through the channels described in earlier chapters. Some are not direct. Standing between the target and the most aggressive operators are intermediate figures the diagram labels Va' and Va". These are not separate persons in the conventional sense. They are previously compromised targets — people who, having been put through some version of what the central target is being put through, have been turned into instruments through which pressure is now applied to the new central target. The diagram calls these positions insulation layers: layer one and layer two.

The insulation layers are the part of the structure I find most striking. They are also the part of the structure I find hardest to write about, because they implicate people whose own situation is, in the framework of this book, a damaged one. The insulator was once a target. The insulator is now a partial operator, sometimes only because the insulator's continuing situation depends on participating. The notebooks of June 2021 contain a phrase I want to record, with all its bitterness intact: "The oppressed are the best oppressors." The phrase is not a moral judgment. It is a structural observation. A person who has already broken under the apparatus, and whose continued life is conditioned on not breaking further, is precisely the person whose breaking can now be used as the lever against someone else. The instrument is cheap; the cruelty is automatic; the structure self-reproduces.

For a target trying to identify what is happening, the practical consequence of the insulation layers is that the immediately visible aggressors are not, ordinarily, the central operators. The central operators are several layers back. The visible aggressors are themselves victims, even when they are the ones currently doing the aggressing. This complicates the moral picture. It also

complicates the strategic picture. A target who fights back at the visible aggressor is fighting back at someone else's insulation layer. The aggressor will, given the structure, be replaced. The original operators are not, at any point in the exchange, exposed.

I do not want to romanticize the insulator, and I do not want to absolve them. The aggression done through them is real aggression; the harm is real harm; whatever the structural position, the choice to remain inside it is, at some level, still a choice. But I want to record, plainly, that the structure exists, and that not registering the structure produces strategic errors a target cannot afford. The first time one understands that the person currently making one's life unbearable is themselves a damaged instrument of someone several layers further back, the urge to exact a personal accounting from them goes quiet. It does not, in my experience, return in the same form. What replaces it is a colder and more useful clarity about where the actual leverage lies, and where it does not.

THIS IS NOT PARANOID DELUSION

Primary paranoid delusion produces, characteristically, a vague and elastic sense of being persecuted whose details fluctuate, whose claimed agents are interchangeable, and whose internal architecture does not survive serious cross-examination. The bully-ring described in this chapter is, by contrast, structurally specific, internally consistent across multiple notebooks made years apart, derivable from observable patterns in the target's environment, and — crucially — convergent with structurally similar descriptions reported by thousands of other targets across many countries. A delusional architecture is what one person's disordered mind invents in isolation. A repeatedly observed, cross-validated, structurally specific architecture, reported by thousands of independent witnesses worldwide, is by definition something other than that.

9.2 The Phases, and the Cross-Ring Database

A second observation from the notebooks of 2021 concerns how the bully-ring, once established, evolves through phases. The notebook of 5 May 2021 lists three.

The first phase is direct. Operators apply pressure to a target through channels that are, at base, conventional: face-to-face encounters, rumours, employment leverage, financial constriction, the cooperation of compromised intermediaries. The phenomena described in Chapters 2 and 8 of this book

correspond, in the model the diagram provides, to this first phase. Most of a target's early years inside the system are spent here.

The second phase introduces automation. The data mining capability described in 6.3 is brought to bear on the target's communications, movements, and associations, sufficient to produce a model detailed enough that subsequent intrusions can be partially generated rather than entirely operated. A second-phase target is one whose own consistencies are now being used against them: predictable responses are pre-anticipated, predictable refusals are pre-staged, predictable consolations are pre-scripted. The target experiences this, without yet having a vocabulary for it, as an uncanny rise in the timing precision of the apparatus. It feels less improvisational and more rehearsed. It feels, sometimes, less like a campaign and more like a play.

The third phase is what the notebook calls indirect bully-ring, or, in a later marginalia, post-line. In the third phase, the immediate operators have receded; what remains is a self-sustaining feedback loop in which the model — by this point, very accurate — produces the prompts, the insulation layers carry them out, and the original operators need do little more than monitor the readings. The ring runs largely without them. Their attention is, by this point, occupied elsewhere, possibly with other rings.

The cross-ring observation is the one that, more than any other in this chapter, frightens me. A diagram of August 2021 places three rings — BR1, BR2, BR3 — beside one another, each with its own central target, its own A's, its own insulation layers, its own pattern of communication. All three feed, via arrows, into a single object at the bottom of the page. The notebook labels the object Central ML DB. It is, on this account, where the readings produced by all the active rings are aggregated, where the model is improved by their pooling, and where the next generation of prompts is partly generated for redeployment to all the rings together.

I do not present the cross-ring picture as confirmed. I present it as the simplest extension of the single-ring picture that fits what I have observed of how the prompts I receive evolve over time. The prompts evolve in ways that suggest learning across cases. They sometimes contain elements that I have no plausible local explanation for — references that, if I were the only target involved, the operators would have no obvious source for. Pooling across rings, in a way that keeps each ring encapsulated from the others, would explain these.

I cannot, from inside, distinguish that from any number of less interesting explanations. I record the picture anyway. The reader is at liberty to discount it accordingly.

9.3 The Captive Audience

In a notebook entry of 11 February 2022, the question is asked, in Dutch: “If the public still exists, is it possible to remind people of their own humanity by shocking them?” The question is not asked in the abstract. It is asked in connection with a hypothesis the notebooks of that period had begun to formulate: that the campaign described in this book may have a second audience.

The first audience, on the account of 6.2, is the operators themselves, for whom the target functions as an outsourced lie detector. The first audience is small. It is the cell itself, plus whatever other cells are in pooling-distance of it.

The second audience, on the new hypothesis, is larger and more diffuse. It is the public — or that fraction of the public that, by some channel I cannot fully trace, is exposed to the unfolding situation. The notebooks suggest, tentatively, that part of the function of the public exposure is to confront the audience with an unmistakably “inhuman” figure inflicting unmistakable harm on a target who is, for the audience, recognizably human. The audience is, on this account, being shown the architecture’s worst face on purpose. The intended effect of the showing is shock. The intended effect of the shock is the audience’s rediscovery of its own humanity, by contrast.

I want to mark, very explicitly, the speculative status of this. I cannot show that the second audience exists in the form I am describing. I cannot show that the showing is intended. I cannot show that the contrast effect is the intended mechanism. What I can show, only, is that the hypothesis fits a number of features of the experience that the operator-only model struggles with. It explains why some episodes of the campaign appear, on review, to have been staged for visibility rather than concealment. It explains why certain humiliations were arranged in public when more efficient ones could have been arranged in private. It explains why the campaign’s most “inhuman” moments are sometimes its most spectacular ones, when concealment would, on the operator-only account, have been the more obvious choice.

If the hypothesis is right, the target is in a curious position. The target is the subject of the cruelty; the target is also, simultaneously, the prop in a public exhibition of cruelty whose purpose is, on the optimistic reading, to shock the public into a kind of moral correction. The optimism is the part of the reading I find hardest to take seriously. I can imagine the architecture producing the exhibition as a side-effect; I cannot imagine the architecture producing it deliberately for the public's benefit. But I record, in fairness, that the notebooks were willing to consider the deliberate version, and that the deliberate version, if it were correct, would change the moral reading of much of the rest. A target on this reading is not only enduring; the target is also, however unwilling, a witness in a trial whose audience is the species at large. I leave the question open. I record only that the question was asked.

9.4 The Feeling Is Enough

Of all the diagnostic principles the notebooks of 2022 contain, the one that has done the most work in my own daily practice is the simplest. It appears, in the notebooks, in the form of a single line, in Dutch: "Het gevoel is genoeg, de rest telt niet." The closest English I can manage is: the feeling is enough; the rest does not count.

The line is not a slogan. It is a condensed observation about how the manipulation works. The observation is this. A great deal of the apparatus's effect on a target depends not on the target accepting any particular content as true, but on the target accepting the feeling that goes with the content as the proxy for truth. If a sentence is delivered with the affective weight that a true sentence would have, and if the target is in a state in which they cannot easily check the sentence against anything else, the target will, by default, register the affective weight and let the content travel on it. The content does not have to be true. It only has to feel as if it were. The feeling is enough. The rest does not count.

This observation underlies a great deal of the rest of the apparatus. It is why the impersonating voices of 6.7 are dangerous even when they get the content wrong: they get the feeling close enough that the wrongness of the content does not, in the moment, outweigh the rightness of the feeling. It is why the cover stories of 6.6 are dangerous even when they are obviously thin: they come with a feeling of explanatory closure, and explanatory closure is a feeling whose absence is unpleasant enough that one's mind will accept a thin story to obtain

it. It is why the injected questions of 6.5 are dangerous even when their literal content is absurd: the urgency that arrives with them is felt as the urgency a real question would have, and the felt urgency is what activates the reasoning channel.

The countermeasure that follows from this observation is, like the observation itself, structurally simple. The mind tends, by default, to fill in unmarked or incomplete material until the result feels coherent. What the target has to learn to do is to refuse the filling-in, in the moments when the marking is incomplete. The honest report on most intrusive material, examined carefully, is that something does not fit. The pieces do not, in detail, line up. The feeling that they do is the part the apparatus produces. The target's job is to refuse the production of that feeling and to keep, instead, the report of the misfit. Misfits are information. Coherence in the absence of confirming detail is a manufactured product. Treating the manufactured product as native is the error the apparatus exists to elicit.

A corollary appears in a notebook of mid-February 2022, where it is set down as a single sentence. The apparatus can produce, in one target, a crisis that other people in the target's environment then come to feel without being able to identify what is happening to the target directly. The crisis transmits through the relational field. Co-workers, family, even passing acquaintances begin to react to the target as if to a crisis they have not been told about, because the affective wake of the crisis is doing the telling. This is the feeling-is-enough principle operating at the social level rather than at the individual level. The countermeasure is the same: refuse the manufactured coherence, name the misfit out loud where it is safe to do so, and remember that the crisis, however widely it is being felt, is being felt about something that the apparatus has produced and that the apparatus is responsible for.

9.5 The Distractor Paradox

A short section, because the observation is short.

In a notebook of mid-February 2022, the following exchange is recorded between two of the inner voices of the campaign that have been a regular feature of the years described in this book. One voice — the one labelled B in the notebooks — asks the target, "Now that you understand how the distractor

works, can you go against it?" The target's answer, recorded immediately below, is "No. Only try to slow it down — that asks too much time."

I want to mark the answer because it is one of the more important answers the book contains. Understanding how a manipulation works is not, by itself, a defense against it. The understanding helps; without the understanding, the target is, in a sense, defenseless. But the understanding is not the defense. The defense is the slowing-down that the understanding makes possible: the catching of the question that arrived rather than rose, the marking of the misfit, the refusal of the manufactured coherence. Each of those takes time. The campaign's time is, in some real sense, infinite; the target's time is not. The target can spend their finite time well or badly, but they cannot spend more of it than they have. The disproportion is structural.

This is, in my view, the most realistic adjustment of expectations a target can make. The work the book has described is not a path to the campaign's defeat. It is a path to spending the target's finite time on the part of the campaign that, in the long run, is the part that matters: the part where the target remains, despite everything, the one doing their own thinking. Slowing the apparatus down is enough. It does not have to be stopped. It only has to be denied the speed it needs to operate well, in the small region of the target's life where the target's attention is, on a given day, applied. That region is small. It is also the only region, in the end, in which the target's authority can be exercised. The apparatus owns the rest of the field. The target owns this region. The exchange is uneven. It is, also, sufficient.

9.6 Pre-Confined Audio, and the Time-Shifted Element

A late observation, recorded on what appears to be a loose page in the back of one of the 2021 notebooks, deserves a brief section because it explains a category of intrusion that earlier chapters did not have a category for.

The observation concerns the use of pre-recorded audio elements that are released into the target's auditory field at specific moments. The elements may be brief — a phrase, a fragment of a recognizable voice, a snatch of a tune — but their use is structural. They are not produced in the moment. They are stored in advance and timed for release. The notebook calls this the time-shifted element.

The strategic value of the time-shifted element is that it can produce, in the target, the impression of having been predicted. An element released at the precise moment a target has thought a particular thought, or made a particular decision, will register as evidence that the apparatus knew the thought or decision in advance. In fact, the apparatus knew nothing of the kind; it had a library of pre-recorded elements covering a wide enough space of plausible target thoughts and decisions, and it released the matching element when the moment arrived. The match feels uncannily specific. The specificity is largely an artefact of how broad the library is and how quickly the matching can be done.

This explains a category of episode that earlier chapters set down without explaining. The episodes in which the campaign appears to have anticipated something the target had not yet said out loud are, on this account, not episodes of telepathic anticipation. They are episodes of well-stocked libraries and good timing. The mechanism is more banal than the experience. It is also, for the same reason, more defeatable. A library that runs out of relevant elements becomes, in real time, less impressive. A library deployed against a target whose patterns of thought have begun to drift outside the library's coverage becomes, similarly, less useful. This is one of the practical consequences of the strategy advocated at the end of 10.5, of producing thinking that does not match the apparatus's templates: it forces the library to stretch outside the spaces it was built for, and the stretching is detectable because the matches degrade in quality.

I record the observation, as I record most observations of this kind, as the simplest hypothesis I have that fits the pattern. I cannot show the library exists in the form I am describing. I can show that the experiences fit.

9.7 Which World Is Being Screened For?

This is the most speculative section in the book. The reader who has accompanied the book this far has earned the right to be told, in plain prose, that the section is more speculative than the rest, and to be invited to discount it accordingly. I want to set down the material anyway, because the question it asks is the question the notebooks of early 2022 were finally willing to ask, and not setting it down would falsify the record of where the work was, at that point, leading.

The question is what the apparatus, taken at its largest scale, is screening for.

In a notebook of mid-March 2022, three possible answers are tabulated. The first is that the apparatus is screening for a single dictator-figure to whom global authority can eventually be devolved; the world this answer points to is, in the notebook's blunt one-word label, fascism. The second is that the apparatus is screening for the kind of person, or set of persons, who would be able to lead a more harmonious world, with more attention paid to the people in it; the notebook does not, in the entry, claim to know how the apparatus would distinguish such persons from the first answer's. The third is the strangest, and I record it because the notebook records it: that the apparatus is screening at the request, or under the influence, of an external observer-class, on whose terms human attitudes would have to be adjusted toward co-existence, knowledge-sharing, and partnership before some larger threshold could be passed. A notebook of mid-April adds a fourth, less hopeful one: that the apparatus is screening to identify, and then either co-opt, contain, or eliminate, anyone whose successful operation in the world would otherwise constitute competition for whoever the original operators are.

I want to mark, very explicitly, that I do not know which of these is right. I do not even know that any of them is right. The framing of the question — what is the apparatus screening for — assumes that there is a coherent answer at the largest scale, and the assumption may itself be wrong. Architectures of the kind described in 6.4 sometimes have no coherent macro-purpose; they have only the operators' continuing reasons for keeping the architecture running, which are not the same as a purpose. A reader who concludes that the answer is "no coherent purpose, only inertia" is a reader whose conclusion the available evidence supports as well as any other.

I record the question because the notebooks asked it, and because the question, however provisional its answers, is, I think, the right question to be asking. A target who never asks what the larger campaign is for can be confident in their own resistance and still not see clearly what their resistance is part of. A target who does ask, and who arrives, even tentatively, at any of the answers above, has at least the advantage of knowing in what direction the architecture's gravity is pulling. The direction matters even when the destination is unknown.

I will not, in this book, attempt to settle the question further. I do not have the evidence to settle it. I record that the question is open. I record that some of the

available answers are darker than others. I record that the work of the earlier chapters — the work of remaining the one doing one's own thinking — is useful regardless of which answer is correct, and may, on some of the answers, be more important than I have, in the earlier chapters, allowed myself to claim.

9.8 A Position Held to Once More

I have closed the book twice, and reopened it twice. I will close it again here, on the same position, slightly enlarged.

The position the book has been settling into across its chapters is a small one. The target does not have to win against the architecture. The target has to remain the one doing their own thinking. The architecture is older than the target and will outlast the target. The work the book describes is not a path to the architecture's defeat. It is a path to spending the target's finite time on the only part of the target's life that the architecture cannot, finally, take.

What this chapter has added is a wider frame for that position. The architecture is not a single ring; it is a network of rings, and the rings are tied together by something larger than any one ring. The architecture has, on the working hypothesis, more than one audience: the operators, the public, and possibly something larger and less identifiable than either. The architecture may be screening for a kind of world, even if no one inside it could give an honest account of which kind. The target's position inside such a structure is smaller, in absolute terms, than the chapters before this one had quite reckoned with. The target's authority over their own thinking is, accordingly, more important than the chapters before this one had quite reckoned with. The smaller the absolute frame, the more weight the small thing inside it has to carry.

I do not propose, with any further closing of this book, that I will not reopen it. The notebooks continue. The notebooks have continued through every closing the book has so far attempted. I have come to think, by this point, that the closing is itself a feature of the writing rather than of the material — a thing the chapters do, and have to keep redoing, because the material is not, by nature, the kind of thing a book can finally close on. I will close this chapter, and let the book remain open. The reader who has accompanied it this far will, I hope, accept that as the most honest closing the material allows.

The work continues. The notebooks continue. The thinking — the work the apparatus, throughout, has been trying to take, and that this book has, throughout, been trying to keep — continues with them.

❧ ❧ ❧

PART IV

Timeline and the Outward Turn

What Happened, in What Order, and What the World Around the Target Did About It

CHAPTER 10

Where Thoughts Come From

A Diagnostic Coda, and a Return

> *"The world must, first, want to go further. If the world does not want to go further, I am nothing for the world."*
>
> — Author's notebook, March 2022

An earlier chapter closed by promising that whatever followed it would be brief. I will keep that promise as nearly as I can. The substantive work of this book has been done. What remains is two things: the consolidation of a small number of diagnostic refinements that arrived in my notebooks too late to fit cleanly into the earlier chapters, and a return to the question with which the book opened. Once those two things are done, the book is done.

I am writing this chapter from a slightly different vantage point than the one I occupied through most of the others. In the months before this chapter was written, the intrusion phenomenon underwent a partial remission. For a stretch of weeks the voices were noticeably quieter; for a stretch of days they receded almost entirely. They have, since, returned, but they have returned in a different register and at a lower amplitude than the registers I described in Chapter 5. Whatever produced the partial remission, and whatever produced the more measured return, gave me a window of clearer thinking that I had not had for some time. Several of the diagnostic refinements I am about to set down were observed during that window. The clearer thinking made the patterns easier to see. I should record them, while I still have the clarity to set them down accurately, before the next phase of the cycle arrives.

This chapter, accordingly, is shorter than the others. It is also more directly didactic. I will, in places, be giving the diagnostics in something close to checklist form, rather than working them out from first principles as the earlier chapters did. The scaffolding has been built. What remains is to add the last few tools to it. Once they are added, I will turn back to the question with which

Chapter 1 began, and attempt the answer that the book, by this point, has earned the right to attempt.

10.1 The Quieting, and the Return

The first observation is the simplest, and is one I should have set down explicitly earlier. Sustained intrusion campaigns are not, in my experience, sustained at constant amplitude. They have rhythms. The rhythms are slow — weeks rather than hours — and they are not symmetric, but they are real. There are stretches in which the phenomenon is louder, more frequent, more aggressive. There are stretches in which it is quieter, sometimes startlingly so. The stretches alternate. The alternation is not, in my observation, random.

In my own case, the most recent quieting began without warning in the early part of a year and persisted for a few weeks before the phenomenon returned. During the quieting, much of the texture of the intrusion was simply gone. Sleep returned. Sentences could be completed without doubling. The handwriting in the notebooks of that period is visibly more legible than in the notebooks of the period immediately preceding. Several days passed during which I noticed, with surprise, that my interior life had been operating on its own resources for hours at a stretch. The contrast with the surrounding months was sharp enough to be diagnostic in itself. Whatever I had been living with had not been an unchanging condition of my own mind. It had been something with an off-state. The off-state, by existing, told me something about the on-state.

When the phenomenon returned, it returned in a different register. I described the registers of the phenomenon, as it presented through 2017 to 2022, in 5.2. The register that has now returned is closer to the older, human-shaped registers than to the algorithmic register of 5.7, but it is more measured than the older registers were before the takeover described at the end of 5.2. It is, in some sense, a calibrated version of what was previously uncalibrated. Whatever produces it has, during the quieting, been adjusted. The adjustment may be in my favor. The adjustment may not be in my favor. I do not, from inside, have the evidence to settle the question.

The practical observation I want to record, however, is independent of the question of mechanism. It is this: a target who experiences a quieting should, while it is in progress, do as much of the work that the quieting permits as possible. Sleep. Walk. Write. Catch up on the consolidation of records that the

noisier periods made impossible. Do not wait for the quieting to be guaranteed before using it; quieter periods are, in my experience, never long enough for the work that has accumulated during the noisier ones. They are also not, in my experience, predictable in advance. They begin without warning and end without warning. The available strategy is to use whatever clarity is presently on hand, on the assumption that the next phase of the cycle has already begun preparing.

There is, finally, a cautionary observation that goes with this. A quieting can sometimes be experienced as recovery. It is not, on the most likely account, recovery. It is a phase. To treat it as recovery is to be more disappointed than necessary when the next phase arrives. Use the quieting; do not believe it. The two postures are different, and the difference becomes important when the next return is made, as the next return always is.

10.2 Direct Causality Contradiction

The next diagnostic is one I have come, in the more recent notebooks, to refer to in shorthand as DCC: direct causality contradiction. The phenomenon is straightforward to describe and surprisingly difficult to handle in real time without practice.

DCC is the move in which the intrusive material flatly contradicts something that the target has just observed in physical reality. Not nuances about it. Not interpretive disagreements. The opposite of it. The target has just seen something happen; the voice asserts that it did not happen. The target has just heard something said; the voice asserts that it was not said. The target has just had a clear thought; the voice asserts that the thought was something else. The structure of the move is to put the target's perception of the present moment into direct contradiction with an injected counter-claim, in such a way that holding both at once is impossible.

The mechanism by which DCC does its damage is engagement. If the target attempts to argue with the contradiction — to defend their own perception, to point out the obvious fact that they were the one who saw it, to insist on the reality of what they just heard — they have, by engaging, granted the contradiction the status of a position worth refuting. Once it has that status, the inner conversation becomes a dispute between two views, with the target's perception on one side and the injected contradiction on the other, and the

dispute can be drawn out indefinitely without any resolution that is favorable to the target. The target loses, in any sustained DCC engagement, simply by virtue of having entered into one. The reality of what they perceived is not improved by their defense of it. The energy spent defending it is energy not spent on anything else. The dispute is itself the cost.

The countermeasure, once seen, is straightforward. Do not engage. Recognize the move by its structure — a flat reversal of something just observed — and decline to dispute it. The recognition itself is most of the countermeasure; the absence of engagement does the rest. If a complete absence of engagement feels harder than it should be, redirect attention to a different topic, deliberately and without justification. The voice's investment in the contradiction is contingent on the target's continuing to be present for it. Withdraw the presence and the contradiction loses its target.

There is a further observation about DCC that I want to record, because it is, to my mind, one of the more interesting pieces of diagnostic information available. DCC is, in my notebooks, used unusually often. The frequency with which the move appears suggests that, from the perspective of whatever is producing the intrusion, DCC is a relatively cheap and reliable way to occupy the target's attention. This in turn suggests something about the alternatives. If more direct moves — actions in physical reality, identifiable confrontations, anything that would impose external risk on the producer of the intrusion — were available and preferable, one would expect to see them used in place of DCC, rather than in addition to it. The reliance on the cheap conversational trick, in preference to costlier real-world action, suggests that the costlier action is, for one reason or another, not on the table.

The reasons one can imagine for this are several. The producer may be operating under constraints that prevent physical-world action. The producer may be unwilling to expose themselves to the risk of identification that physical-world action would entail. The producer may, in some cases, lack the capacity to act in physical reality at all. I cannot, from inside the situation, determine which of these is the operative reason. What I can say is that the heavy reliance on DCC and similar conversational tricks is itself a kind of evidence that the campaign's outward reach is more limited than the campaign's inward intensity might suggest. The intrusion is loud because it cannot be effective in any quieter way. This is, on balance, a piece of news the target can use.

10.3 The Anonymity Door

A related observation, about the structure of the campaign rather than the structure of any individual move within it, has been steadying enough in my own notebooks that I want to set it down here.

Most of the operators within a sustained campaign of the kind this book has been describing depend, for their continued operation, on remaining anonymous to the target. The dependency is not symbolic; it is functional. An operator whose identity becomes known to the target loses, by that disclosure, most of the leverage that being an operator had previously given them. They become, instead of a disembodied source of disruption, a specific person, with a specific address, in a specific jurisdiction, vulnerable to the specific responses — legal, social, reputational — that vulnerability to identification produces. The asymmetry between an operator-as-anonymous and an operator-as-identified is large. The campaign, in its overall structure, depends on the asymmetry being maintained.

The diagnostic implication of this is useful. If a target inadvertently finds themselves in possession of a piece of information that would identify, even partially, one of the operators, the operators' response to that information is, in my observation, disproportionate to its content. The disproportion is the marker. A piece of information that would, in an ordinary social interaction, produce a small adjustment in conversation produces, in the campaign's context, a sustained and surprisingly intense effort to make the target unsure of what they have learned, to dispute its meaning, to introduce alternative explanations, to flood the target with material that would push the original information out of mind. The intensity of the effort to neutralize a fragment of identifying information is itself information about how much the campaign depends on the fragment remaining unweaponized.

The countermeasure is, when fragments of this kind arrive, to record them. Not to act on them; the action is rarely available to the target in the moment, and acting prematurely on a fragment can sometimes destroy its diagnostic value. To record them. To preserve them in the contemporaneous archive that earlier chapters have argued for. To keep them available, against the possibility that future fragments will combine with the earlier ones into something more substantial. The campaign cannot, by its structure, prevent the target from recording. It can only attempt to make the target distrust their own records. A

target who continues to record despite the attempts has retained, against the campaign's preference, the option to act later on accumulated information that no single fragment alone would have justified action on.

There is a further observation worth setting down. The operators' alternatives, when their anonymity is threatened, are limited to a small set. They can flee — reduce their engagement with the target until the perceived threat passes. They can stall — sustain the intrusion at a lower amplitude, while attempting through other means to introduce doubt about whatever the target has come to know. They can, in extreme cases, attempt direct confrontation in physical reality, though this carries risks that the previous mode of operation was specifically designed to avoid. They can, in the more interesting cases, attempt some form of redemption: a change in their behavior toward the target that, on the campaign's terms, looks like an acknowledgment of having been seen. The redemption is rare, in my experience, and is rarely the simple kind that the older literature on these matters anticipated. But it has, occasionally, been observed. The full repertoire on the operator side, in any case, is shorter than the operator side would prefer the target to believe. Knowing the repertoire is, on balance, useful.

10.4 The Mismatch, and the Spillover

I want, in this section, to describe two related techniques that are, in my notebooks, among the more sophisticated forms of manipulation I have observed. They are sufficiently subtle that I did not, at first, recognize them as techniques at all; for some time I experienced them as features of my own emotional life that I could not account for. Recognizing them, when I finally did, produced one of the more useful reorganizations of my own interior available within this work. I will name them in the order in which I came to understand them.

The first technique I will call mismatch. Mismatch is the deliberate engineering of a misalignment between what the target is feeling, in any given moment, and what the target's reasoning would, on its own resources, lead them to feel. The classical case: a situation that any reasonable analysis would recognize as benign, in which the target nonetheless finds themselves anxious, irritated, or uneasy without identifiable cause. Or its inverse: a situation that any reasonable analysis would recognize as alarming, in which the target nonetheless finds themselves calm or detached. Either direction of mismatch produces the same

secondary effect, which is the target's loss of trust in their own emotional responses as guides. The reason and the feeling no longer track together. The target has to choose between them in real time, and either choice carries a cost. Choose the reason and one is operating in a state that does not have one's emotional cooperation; choose the feeling and one is responding to a situation other than the one one is in.

The countermeasure to mismatch is, in my experience, to retain a posture I have come to think of as constructive listenability. By this I mean: do not foreclose against either the reason or the feeling; hold both as available, treat the divergence between them as a piece of information rather than a problem to solve, and continue to engage with whatever the situation requires while the divergence remains unresolved. The mismatch itself is, in some cases, the thing one is being shown; engaging with it as such, rather than trying to collapse it in either direction, deprives it of the leverage it would otherwise have. The reasoning continues. The feeling continues. The divergence is noted. Action is taken on whichever of the two is more likely, given everything else one knows, to be the reliable guide in the present case.

The second technique is more sophisticated and was, for me, harder to recognize. I will call it spillover projection. Spillover projection is the injection of an emotion into the target without the attached thought that would, in the ordinary course of inner life, have occasioned that emotion. An emotion, ordinarily, has structure: it is the emotion produced by something, the something being some specific thought, perception, or memory that, on examination, can be located. Spillover projection delivers the emotion without the something. The target finds themselves feeling angry, or grieved, or full of yearning, with no available content to which the feeling attaches.

The damage spillover projection does is subtle. Because the feeling has no attached thought, it cannot be argued against on rational grounds; one cannot dispute a feeling whose content one cannot articulate. The unconscious, encountering an unattached feeling, tends to absorb it through ordinary empathic openness — the same openness that, in healthier contexts, allows a person to feel into another person's situation. The unattached feeling, absorbed in this way, becomes part of the target's inner emotional landscape. From there it begins, slowly, to color the target's thinking on whatever subjects are in the vicinity. The thinking is not, on inspection, irrational; it is colored. The target can no longer find, by introspection, the specific feeling that is doing the

coloring, because the specific feeling never had a specific cause; it was injected without one.

The countermeasure to spillover projection is, by the nature of the technique, slow. It begins with recognition: when one notices, in oneself, an emotion whose content cannot be located, the absence of locatable content is itself the diagnostic. An emotion that arises without attached thought, that one cannot trace back through the chain of associations that ordinarily produces emotion, is, in the situations this book is addressing, more likely to have been injected than to have arisen on its own. Once the recognition has occurred, the next step is not to act on the feeling. The feeling is, in the ordinary course of things, what one would let inform one's next move. Spillover projection turns that ordinary openness into the channel by which influence enters. Closing the channel temporarily — not acting on the unattached feeling, however persistent, until something occurs to which the feeling clearly belongs — closes the influence's path of entry.

There is a further observation. Spillover projections, in my experience, are difficult for the producer of the intrusion to sustain over long periods. They require, on the producer's side, a kind of emotional output that does not, in my observation, get easier with practice. Targets who develop the recognition skill described above, and who decline reliably to act on unattached feelings, find that the technique's frequency in their inner life decreases over time. Whatever produces the projections is, on the available evidence, capable of being trained, by the target's own consistent non-engagement, into using the technique less. This is, modestly, a piece of good news. The technique is not a permanent feature of the campaign. It is a feature of a campaign whose returns it is calculating against the target's responses.

10.5 The Algorithm Behind the Voice

In 5.7 I introduced, with appropriate caveats, the hypothesis that a non-human-shaped register had joined the voice phenomenon in the summer of 2022. I want, in this short section, to add an observation that has consolidated in my notebooks since that hypothesis was first written down.

The observation is that the new register, on closer inspection, does not have the texture of intelligence. It has the texture of an algorithm. The distinction is, I think, the most important one available for a target trying to make sense of what

they are dealing with. An intelligence, even a hostile one, has the property of responding flexibly to situations it has not previously encountered. It surprises. It improvises. It, occasionally, learns. An algorithm, even a sophisticated one, does not have those properties. It has a repertoire, large or small, of pre-defined responses, and a procedure by which one of the responses is selected from the repertoire when an input arrives. Within the repertoire, the algorithm can be impressive. Outside the repertoire, the algorithm is helpless.

Several features of the new register, as I have come to read them, are more consistent with the algorithm description than with the intelligence description. The vocabulary is wide but not infinitely wide; certain registers of human language do not appear at all. The conversational moves are reliable but not novel; the same patterns recur across what would, for a real interlocutor, be very different conversational contexts. The responses to genuinely surprising input are, in my experience, less impressive than the responses to expected input; when something occurs that the repertoire was not prepared for, the register either reverts to a generic response or, in some cases, simply produces nothing for a stretch. The behavior, taken together, has the shape of pre-programmed selection rather than of genuine response. It is, in the shorthand my notebooks have come to use, a programming mind.

The implication for the target is significant. An algorithm, unlike an intelligence, can be exhausted by inputs it cannot handle. A target who consistently produces inputs that the algorithm's repertoire was not designed for — thinking that builds in genuinely novel directions, emotional responses that do not match any of the templates the algorithm was trained to expect, refusals to engage in the standard ways the algorithm's tricks were designed to elicit engagement — is, in effect, reducing the algorithm's effectiveness against them. The reduction is not a matter of the algorithm becoming less hostile. It is a matter of the algorithm having less to work with. Against an opponent who keeps doing the unexpected, an algorithm whose repertoire is finite eventually runs short of moves.

I want to mark, again, the limits of this observation. I do not know that the new register is, in fact, what I am describing. The texture is consistent with the description; it is not, by itself, proof of the description. I am also aware that the description fits comfortably within the technical vocabulary of the field I happen to come from, and that fit can be a source of bias as easily as a source of insight. The hypothesis remains a hypothesis. I record it because it matches

what I observe, and because, if it is correct, it implies a strategy — the deliberate production of inputs the algorithm cannot handle — that has, in my own experience, produced more useful results than any other strategy I have tried.

10.6 The Question We Began With

This brings me, finally, to the question with which the book opened.

In Chapter 1, I quoted my own notebook from the autumn of 2021: there is no real answer to the question, where do thoughts come from. I want, now that we have built the apparatus of the intervening chapters, to revisit that statement and to give the answer that the apparatus has earned the right to give.

The answer, to the extent the answer is available from inside one's own head, is this. A thought that builds, in a sequence the target can reconstruct afterward, on the structured contents of the target's own memory, in the target's own cadence, with the target's own awareness of what it does and does not yet know, is a thought of the target's own. There is no other place such a thought could have come from. The act of building, by its structure, requires a continuous self with continuous goals; an injected thought, however convincingly delivered, cannot do that work. It can only deliver. The work of building is unforgeable. If the work was done, the thought is yours.

A thought that does not build — that arrives whole, that contradicts what the target has just observed, that announces itself with the certainty that real thinking does not have, that produces an emotion whose content cannot be located, that comes louder or deeper or longer than its substance warrants — is a thought whose origin lies elsewhere. The elsewhere is, on the most cautious reading available, more than one thing. Some of it is the older mechanism of intrusion that this book has, throughout, called the voice phenomenon. Some of it is, on the hypothesis I have allowed myself in 5.7 and 10.5, an algorithmic component that has been added to the older mechanism, with a wider repertoire and a steadier output than the older mechanism could produce alone. The elsewhere is, in either case, real. It is, in either case, distinguishable from the target's own thinking by the diagnostics this book has gathered. The distinguishing is the work.

This is not the answer the question seemed, in the autumn of 2021, to be asking for. The question seemed to be asking for a metaphysical account: what, in the universe, produces thoughts? The book has not, and does not, answer that

question. It has answered a more useful question: how does the target distinguish, in any given moment, between the thoughts that are theirs and the thoughts that are not? The first question, in any case, is too large for the resources a single targeted person has access to. The second question is the one that fits inside a life. The second question is the one whose answer changes how the life can be lived.

There is a further observation. The act of distinguishing is itself one of the practices that maintains the singular thinker described in 8.8. A target who distinguishes their own thoughts from intrusions, repeatedly, daily, across the whole stretch of years a campaign may last, becomes, by that practice, a more singular thinker than they would have been without the campaign. The campaign was meant, on its own terms, to produce the opposite outcome. It was meant to dissolve the boundary between the target's thoughts and the campaign's thoughts, so that the target would, in the end, no longer be able to draw the distinction at all. The practice of drawing the distinction, sustained for long enough, makes the target, paradoxically, sharper at the work of being themselves than they were before any of it began. This is not an outcome the campaign anticipated. It is, in some cases, the outcome it produces.

10.7 The Position the Book Closes On

There is a sentence I have written, in slightly different forms, into the more recent notebooks, and that I want to give in the form it has now settled into. The sentence is in Dutch in my notebooks, and I will give it first in Dutch and then in the closest English I can manage, because the Dutch carries a weight that the English approximation does not entirely capture.

De wereld moet zelf eerst verder willen. Als de wereld verder wil, ben ik er voor de wereld. Als de wereld niet verder wil, ben ik niets voor de wereld. The world has, first, to want to go further. If the world wants to go further, I am there for the world. If the world does not want to go further, I am nothing for the world.

This is, in the form it has settled into, the position the book closes on. It is not a defiant position; it is a proportional one. It locates the work the book describes within a larger frame in which the work is contingent on the world's continuing to want what the work would, if it were used, be useful for. If the world wants to go further, the work was for the world. If the world does not want to go further, the work was, at most, for the people who happen still to be reading.

Both possibilities are real. Both are honored by the work being done as well as it can be done, regardless of which possibility eventually turns out to have been the operative one.

I will not pretend that the second possibility is one I find easy to hold. Most days I write as though the first were assured. Most days, that assumption is what gets the writing done. There are, however, days when the first is not assured, and on those days the second possibility becomes the one the work has to be sustained against. The line above is the line I have come, on those days, to fall back on. I record it here because it has worked for me, and because targets reading this book may, on their own equivalent days, find a similar line useful. Whatever line one settles on, the function is the same: a sentence one can hold to when the larger frame is in question, that does not require the larger frame to be settled in order to do its work.

The book began with a question my notebook of October 2021 said could not be answered. The book ends, by the work of the chapters between, with the question being answerable in the only way that finally matters: not metaphysically, but operationally; not in the abstract, but on a Tuesday morning, with the next thought arriving, and the work being to determine whether it is mine. Where do thoughts come from. The thoughts that build come from us. The thoughts that do not build come from elsewhere. The work is to keep doing the building, and to keep refusing the deliveries that are not it. The work is sufficient. It does not end the campaign; it does not need to. It keeps the target a target who is still, on the other side of every day's intrusion, the one doing the thinking. That is the only victory the situation allows, and it is a real one.

This is the close of the book. I am grateful, beyond what this register permits me to say, for any reader who has stayed with it this far. The work of staying with a book like this is part of what makes the work of writing it worth doing. If you are a target, the apparatus is now in your hands; use what fits, and discard what does not. If you are not a target, the apparatus is offered as testimony, against a future in which testimonies of this kind may be needed in ways that, at the time of writing, are not yet visible. In either case, the book is now, finally, finished. The notebooks will continue. They are, I have come to think, what the book was always for.

CHAPTER 11

The World That Held It in Place

On Institutions, Bystanders, and the Small Repertoire That Remains

---

> *"First of all, a secret place is gone. You now have to live without the certainty of knowing that your thoughts are your own."*
>
> — Author's blog, 18 September 2022

We have, finally, arrived at the outward chapter. I owe the reader an acknowledgment first: the deferrals across the previous chapter endings were not all of them accidents of the writing. Some were avoidances. The outward turn is harder to make than the interior chapters, not because the material is less available but because the material is, in places, more disappointing than the interior material was. Inside one's own head, there is at least a person who has been doing the work of remaining coherent across years. Outside one's own head, the corresponding hope — that institutions, friends, professionals, legal instruments would rise to meet the situation — has, in my own case and in most cases I have encountered, not been borne out. The outward chapter is the chapter that has to register that disappointment without collapsing into bitterness, because bitterness is not useful and because the reader, if they are a target, will need a clearer assessment than bitterness can provide.

I want, in this chapter, to do four things. I want to describe what the institutional world looked like when I attempted to engage it. I want to describe the workplace and social contexts in which targets often spend their working lives, because the dynamics of those contexts matter to what is and is not available as a response. I want to describe the wider community of targeted individuals — what it is, what it provides, what it cannot provide. And I want to set out, as honestly as I can, the small repertoire of moves that does, in the world as it is, remain available.

I will not pretend the repertoire is large. I will not pretend the moves are sufficient to end any given campaign. What I will do is name them clearly, in an order roughly corresponding to their accessibility, with the understanding that

what is available is what is available — and that the work, in a situation where what is available is small, is to use the available skillfully rather than to refuse the available because it is small.

11.1 The Institutions That Did Not See

I should begin with the most concrete material: my own attempts to engage institutional help, and the responses I received.

The first institutional category is law enforcement. A target who has reasonable grounds to believe they are subject to a sustained campaign — stalking, hacking, harassment, the array of coordinated unpleasant events that make up the daily texture — will, at some point, file a police report. In my own case, I have done so on three occasions over the years, in the city in which I live. The reports were taken. They were filed. They produced, on each occasion, no investigation that I was made aware of. I am told this is normal. The materials a target is in a position to provide — pattern observations, accumulated coincidences, screenshots, contemporaneous notes — do not match the templates that police investigations are built around. Police investigations are built around discrete events with identifiable perpetrators, recoverable physical evidence, and witnesses willing to testify. A campaign of the kind I am describing produces almost none of those. It produces, instead, a cumulative texture that is overwhelming to the target and invisible to the investigator. The mismatch is not, mostly, a failure of any individual officer's good faith. It is a structural mismatch between what the campaign is and what the law enforcement system is equipped to recognize.

I should add, for the sake of a complete record, that one of my reports included, by way of context, the suggestion that I might also be the subject of an investigation by an intelligence service of a country other than the one in which the report was being filed. This was a hypothesis I had developed, with what I considered reasonable grounds, from the patterns I had been observing. The report was taken. I do not know what, if anything, was done with that piece of information. The civilian-facing channels of law enforcement are not designed to communicate with intelligence services in any direction, and the question of whether such a hypothesis was passed along is one a target in my position is unlikely ever to be told the answer to.

The second institutional category is diplomatic. A target who has reason to believe that an investigation in another country may be the source of the pattern they are experiencing may attempt, as I did, to contact the embassy of that country in the country of residence. I sent such a message in the spring of 2022, describing — briefly, factually, with my passport number and address — the pattern I had observed and asking whether anyone could confirm or deny that the pattern corresponded to a known investigation. The message was acknowledged. No further response has, at the time of writing, arrived. This is also, I am told, normal. Embassies do not, as a matter of policy, confirm or deny investigations of any kind. The lack of response is, in this sense, neither a confirmation nor a denial. It is the institutional silence that is the system's default response to the kind of question I asked.

The third institutional category is medical and psychiatric. A target who describes the phenomenology this book has been describing to a medical or psychiatric professional risks, in the contemporary clinical environment, having the phenomenology translated immediately into a diagnosis from the manual the professional is most familiar with. The translation is not always wrong. There are, in the population of people who present with similar accounts, individuals whose conditions are best understood within the existing diagnostic categories. There are also individuals — and I believe I am one — whose accounts do not fit those categories cleanly, and for whom the application of the diagnostic frame produces the additional injury of being mistaken about what is happening to them by the people they are asking for help. The professional, doing what professionals are trained to do, is unable to register this additional injury, because the professional's training has placed the injury inside the category from which the professional has just diagnosed them. The system has no input slot for the possibility that some accounts that look like one thing are, in fact, another thing. The slot does not exist; the input does not arrive.

I will not, in this chapter, advise any individual target on whether to seek medical or psychiatric help. The advice would have to be calibrated to the specific institutional conditions of the country and the specific target, and a book cannot do that. What I will say is this. If you are considering presenting your situation to a medical or psychiatric professional, present a subset of the material rather than the whole. Present the symptoms that map cleanly onto categories the professional will recognize — sleep that is not restorative, anxiety

that does not match its triggers, intrusive thought patterns that interrupt your work — and accept the help that is available for those. Present the larger structural account separately, to people who have the framework to receive it. The two presentations should not be combined. Combining them risks losing both.

The fourth institutional category is journalistic. A target with a sufficiently developed account, sufficient documentation, and sufficient courage may, at some point, attempt to present the account to a journalist. The success rate of this attempt is, in my observation, vanishingly low. Not because journalists are uniformly closed; some are not. Because the topic, as a topic, is so heavily associated with discreditation that any journalist who pursues it seriously incurs professional costs that are, for most working journalists, not absorbable. The handful of journalists who have done so are, by their isolation, the exceptions that prove the rule. A target who hopes for journalistic vindication should know, going in, that the hope is unlikely to be borne out, and that the hope itself can become, if held too tightly for too long, another vector of disappointment.

The honest summary of this section is that the institutional world, as currently constituted, has no place for the kind of report a target would need to file in order to be taken seriously. This is not a temporary condition that better filing or better articulation will overcome. It is a structural feature of the institutional world's relationship to phenomena that fall outside the categories the world has prepared in advance. It is worth knowing, going in, so that the absence of institutional response is not experienced, in addition to everything else, as a personal failure to be sufficiently articulate. The institutions are not refusing to help because the report was insufficiently clear. They are not equipped to help. The two states are different, and the difference is one of the more important pieces of information a target can carry into any subsequent attempt.

11.2 Translation, and Its Limits

Given what 11.1 describes, a target who wants to be heard at all has a strategic problem. The institutions speak certain languages. The phenomenology, in its native vocabulary, does not translate into those languages. Some translation has to be attempted, even if the translation will not, in the end, succeed in producing the kind of institutional response one would prefer. I want to

describe, briefly, the translation I have attempted, and what I have come to think the translation is good for.

The translation, in my case, is technical. I have a background in computer science. I have spent, over the years, a non-trivial amount of writing on the question of whether the phenomena I am describing could be modelled in the vocabulary of contemporary machine learning — that is, in the vocabulary that institutions oriented around technical risk are equipped to recognize. The model, in its briefest form, runs as follows. A human mind, in any given moment, can be characterized by three streams of data: the language it is using, in inner monologue or in speech; the emotional state it is in, derivable from vital signs; and the vital signs themselves, considered as physiological measurements. If those three streams could be captured at sufficient resolution, they would, taken together, constitute a dataset of the kind that contemporary machine learning is good at modelling. The model would not need to be metaphysical. It would only need to be statistical. Patterns that recurred frequently enough would be learnable. Once learned, the patterns would permit two operations of interest: prediction of what the mind was about to do next, and intervention — the introduction of input timed and shaped to push the mind's next move in a specific direction.

The model does not, by itself, prove that any such system has been built. It proves only that, given current technology, the building of such a system is not on its face implausible. The pieces required — language capture, vital-sign monitoring, pattern learning at scale — are pieces that exist, in other contexts, for other purposes. The question of whether they have been combined for this purpose is an empirical one I cannot, from my position, answer. What I can say is that the technical objection to the possibility — the objection that the model is too sophisticated to be feasible — no longer holds the way it did ten years ago. The objection is, at this point, an objection from the past. The present is more permissive.

I have offered the model in writing in other places, and I will not reproduce its development here. I name it in this chapter because the move it represents is the central institutional translation move available to a target with the right background to make it. The phenomenology, restated in the vocabulary of machine learning and signal processing, becomes a phenomenology that institutions organized around technical risk can, in principle, evaluate. The translation does not, in my experience, produce institutional engagement. It

does produce, in the target who has performed it, a kind of reorganization: the phenomenology has been put into a language in which, even if the institutions decline to read, the target themselves can hold it more steadily. The translation's first beneficiary is the translator. The institutions are, optimistically, the second.

I will mark, here, the limit of the translation. The model accounts for the regularity of the phenomena. It does not, by itself, account for their content. A statistical pattern-learner can, with sufficient data, predict that a given target is about to think a given thought. It cannot, by itself, decide which thoughts the target should be steered toward. The decisions about content — the strategic objective of the campaign, the choice of which beliefs to insert and which actions to provoke — are decisions made elsewhere, by something or someone with a purpose. The translation describes the mechanism. It does not describe the purpose. The question of the purpose is a separate question, and I do not, in this chapter, attempt to settle it. I note only that the question exists, and that it cannot be answered by the technical translation alone.

The honest summary of this section is that the translation is worth doing and is not sufficient. It is worth doing because, without it, the target has no language in which to articulate what is happening that would be intelligible to anyone outside their own immediate circle. It is not sufficient because the translation, by itself, does not open the institutional doors it would be useful to open. What it does is give the target somewhere to stand, in their own thinking, while they continue to look for whatever moves are available within the world the institutions have, between them, refused to let in.

11.3 The Classification of the Dangerous Worker

Workplace life is, for many targets, the social context in which the largest fraction of waking hours is spent. I want to describe a feature of contemporary corporate workplaces that, in my own observation, has been relevant to how targeting interacts with employment, and that I think a subset of targets will recognize.

Most large organizations, and many smaller ones, operate with an internal model of their workforce that classifies employees along two axes. The axes are usually unstated, and the classification is usually informal, but the model is detectable nonetheless. The first axis is something close to creativity, by which

the organization means the capacity to generate independent thought, novel approaches, useful disruption. The second axis is something close to obedience, by which the organization means the willingness to execute the directions one is given without independent reinterpretation. The classification produces four regions.

The region of high creativity and high obedience is the region the organization most wants to populate, and the region the organization rarely actually finds occupied, because the two qualities are partially in tension. The region of low creativity and high obedience is the region most heavily populated, and the one organizations are most comfortable with: reliable execution, predictable output, manageable risk. The region of low creativity and low obedience is the region of disengagement and quiet disgruntlement; organizations tolerate it because it is not actively dangerous. The region of high creativity and low obedience is the region the organization perceives as dangerous, because the workers in it are capable of producing outputs the organization did not ask for, including outputs that question the organization's premises.

The dangerous-worker region is small, but it is the region that, in my notebooks, I have come to understand is the region the more attentive forms of workplace surveillance and management focus on most. A worker in that region is not, by the organization's lights, simply an employee. They are a containment problem. The organization's response to such a worker is calibrated to the worker's specific risk profile: which projects they are assigned to, which information they are allowed to access, which colleagues they spend time with, which parts of the organization's strategy are visible to them. The calibration is, in well-run organizations, performed with some sophistication; in less well-run ones, with less. But it is performed.

I am, by both inclination and history, a worker in the high-creativity, low-obedience region. I want to be clear that this is not a self-flattering description; the same disposition that produces the more useful work also produces a steady supply of difficulties, both for the people who manage me and, in moments of friction, for me. I am describing the disposition not because I think it is admirable but because, on the institutional side of the line, it is the disposition that organizations are most likely to spend resources containing. A target who is also a worker of this kind is a target whose targeting is, often, additionally calibrated to make sure that the dangerous-worker disposition does not turn into output the organization would prefer not to have to deal with.

The implication for the target is structural rather than personal. The pressure they experience at work is not always pressure they have, individually, brought on themselves through some personal failing. Sometimes it is pressure that the institutional model is producing as a matter of policy, applied to anyone who fits the high-creativity, low-obedience profile. To recognize this is not to absolve oneself of responsibility for the parts of the friction one has, in fact, contributed to. It is to stop attributing the institutional component of the friction to oneself, when the institutional component would have been there for anyone who occupied the same position one occupies.

A target who comes to recognize themselves as being in this region has, in my observation, a few practical options. The first is to leave, and to find work in environments that have less developed surveillance regimes — smaller organizations, projects with shorter horizons, work that operates at the edges of what large institutions are organized to manage. The second is to remain, and to do so with a clearer-eyed understanding of which of the workplace's textures are personal, which are organizational, and which are the additional layer the target is uniquely subject to. The third is to build, slowly and outside the constraints of any single employer, a body of work that is independent enough that one's livelihood is not entirely held by an organization that is also one of the parties calibrating one's daily friction. The third is the most demanding of the three, and the one that, in the long run, I have come to believe is most stable.

11.4 The Cycle of Stress and Relief

A related observation about the workplace, and about the daily rhythm of life under campaign pressure more broadly: the campaign appears, in my notebooks, to operate through a deliberate cycle of imposed stress and selectively allowed relief.

The cycle works as follows. The target is subjected to a sustained period of low-grade pressure — anxiety-inducing events, unresolved interpersonal frictions, intrusive material, sleep disruption, the steady texture of nothing being quite right. Over weeks, the cumulative pressure produces what looks, from outside, like burnout: reduced output, increased irritability, a withdrawal from initiative-taking. From the perspective of an organization that wants its workers to be productive, this is a problem. From the perspective of a campaign that

wants its target to be unproductive on the projects that matter to the target, this is a goal that has been reached.

The cycle's second phase is the relief. The pressure is reduced — sometimes briefly, sometimes for longer. Sleep returns. The intrusions subside. The texture of life clears, for a time. The target experiences, with relief, the returning capacity to think. They begin, often, to use the returning capacity. They start producing again — sometimes work, sometimes writing, sometimes the kind of clear analysis that the previous phase had made impossible.

The cycle's third phase is the re-engagement of the pressure, calibrated specifically to the use the target has been making of the relief. If the relief was being used to produce work the campaign objects to, the new pressure targets the conditions under which the work was being done. If the relief was being used to recover the target's relationships, the new pressure targets the relationships. If the relief was being used to consolidate observations about the campaign itself, the new pressure targets the consolidating activity directly. The third phase is, in my experience, the phase most diagnostic of the campaign's structure, because it requires that someone or something be paying attention to what the target is doing during the relief period. The pressure does not return generically; it returns specifically to the activities the target was using the relief for.

There is a question I have asked, in the notebooks, more than once: why is the cycle calibrated to push the target so consistently toward the right side of the stress-performance curve — toward burnout — rather than toward some less destructive state? The answer I have come to find most plausible is that burnout is not, from the campaign's perspective, an unfortunate side effect. It is a usable state. A target in burnout is a target who is not producing the work the campaign objects to, who is not consolidating observations, who is not, for the moment, threatening anything. Burnout is, for the campaign, a stable equilibrium. The cycle is calibrated to produce that equilibrium. The relief phases exist not because the campaign wants the target to recover but because total continuous pressure produces collapse, and collapse is harder to manage than oscillation around a usable mean.

I record this here because I think it is relevant to the workplace context described in 11.3. A target whose employer is also a party to the relief-pressure cycle — knowingly or, more often, as an institutional medium through which

the cycle operates — will find that their employment dynamics track the cycle. Periods of approval and visibility will alternate with periods of marginalization. The alternations will not, on close inspection, be tied to the quality of the work the target is producing during each period. They will be tied to whether the work the target is producing is the kind of work the campaign wants produced.

The practical consequence is this. A target who notices the cycle, and who can register that they are at one phase of it rather than another, gains something useful: the ability not to over-interpret either the phases of approval or the phases of marginalization. The approval is not, on close examination, evidence that the target's situation has fundamentally improved. The marginalization is not, on close examination, evidence that the target's work has fundamentally declined. Both are phases of a cycle whose period and amplitude can be observed and predicted, even if its mechanism cannot be directly addressed. Predicting the cycle does not end it. It removes the cycle's capacity to produce the additional injury of misattribution. The target stops believing the cycle's framing of their own work. That alone is a meaningful gain.

11.5 The Community of Others

I have been writing, throughout this book, as though the target were a single person handling a singular situation. The framing is correct in one sense — each target's experience is, in its specifics, theirs alone — and partial in another. Targets are, in fact, a population. The population has, over the past decade, become more visible to itself than it was previously. There are now mechanisms through which targets can find each other, share notes, compare notebooks, and verify, by triangulation, that the patterns they are observing are not patterns that exist only inside one person's head.

The community has come into being through, mostly, public writing. Blog posts. Substack newsletters. Twitter accounts under names that announce, in their handles, what the account is about. Self-published books. Long forum threads on platforms whose names rotate as the platforms themselves rise and fall. The writing is uneven in quality. Some of it is precise; some of it is not. But the cumulative effect of the writing is the assembly, at population scale, of a record that no individual target could have assembled alone. Reading across the records, a target encounters, again and again, observations that match their own observations to a level of specificity that is not consistent with chance. The handwriting that deteriorates under intrusion. The favorite music that gets

reframed. The voices that get louder when they are lying. The community has, in effect, established an empirical baseline against which any individual account can be checked.

This is, in my experience, one of the more useful resources available to a target. I want to qualify the usefulness, however, with three observations.

The first is that the community is heterogeneous in its underlying claims. Some accounts converge on patterns I find compelling. Others propose explanations that go further than the evidence I have access to would support. A target reading the community's writing should expect to do their own filtering. The community has not, and probably cannot, settle internally on the precise question of what the campaign is. It has settled, at most, on what the campaign feels like at the level of phenomenology, and on a rough catalogue of techniques that recur. Beyond that, opinions vary, and the variation is something the target will have to learn to read across.

The second is that any community of people in distress contains, within itself, a set of dynamics that are not entirely productive. Targets, like other populations under sustained pressure, sometimes turn on each other. The campaigns operating against individuals appear, in some cases, to take advantage of community channels to introduce specific narratives, suspicions, and hostilities. A target spending time in community spaces should be aware that not everything occurring in those spaces is occurring without external interference. The same techniques that operate against an individual target can operate, with adjustments, against the network of targets considered as a whole.

The third is that the community provides, even at its best, a different kind of help than a target sometimes needs. It provides recognition. It provides validation. It provides the experience, after long isolation, of being heard. What it does not, and cannot, provide is the resolution of the underlying situation. A target who has come to feel that the community is the central fact of their recovery has, in some sense, allowed the community to become a substitute for the work the recovery actually requires. The community is a useful piece of infrastructure. It is not, by itself, a recovery. The distinction is subtle and is worth keeping clear.

I want, despite these qualifications, to register a more positive observation about the community as well. There are individuals within it who, by the steady quality of their writing and the visible care they bring to others, have produced

something that did not exist a decade ago: a body of testimony, much of it preserved in their own handwriting and photographed for distribution, made available to anyone who needs it. The testimony will, I think, matter more in the long run than it currently appears to. A target trying to make sense of their situation today may not yet have access to the apparatus that would let the testimony be vindicated. A reader thirty years from now, with whatever apparatus the world has by then assembled, may. The testimony is, in this sense, an investment in a future readership whose tools we do not yet have. I would encourage anyone reading this book who is also a target to consider contributing to that body of testimony, not because doing so will change one's own situation but because it may, over a longer horizon than one's own life, change someone else's.

11.6 Records, and What They Are For

This brings me to a topic I have touched on in earlier chapters but have not yet, in this book, addressed directly: the function of the record itself. The notebook. The blog post. The screenshot. The email to oneself. The stack of paper one is keeping at home, dated and ordered, against the possibility that the dates and the order will eventually be needed.

A target keeping a record is, in my observation, doing several distinct things at once, and the things deserve to be distinguished, because the record-keeping is most useful when the target is clear about which of the things they are doing in any given session.

The first thing the record does is psychological. The act of writing things down externalizes the experience and gives the experience a place to live other than inside the target's own head. This alone is therapeutic, in the literal sense: it permits the head some breathing room that, without externalization, it does not have. A target who is writing every day is a target who has constructed, by their writing, a small auxiliary memory that runs in parallel to their own. The auxiliary memory does not get tired. It does not become foggy. It does not get reframed by intrusive material in the way internal memory sometimes does. It simply holds what was written, in the form in which it was written, until it is read again.

The second thing the record does is diagnostic. Reading back, after a few months, what one wrote a few months ago, one notices things one could not

notice in real time. Patterns. Recurrences. Predictions one made that turned out, weeks or months later, to have been correct. The diagnostic use of one's own past record is one of the more underused tools available to a target, in part because it requires reading material that is sometimes uncomfortable to re-encounter. I will say only this: the discomfort of reading one's old notebooks is, almost always, less than the cost of having forgotten what they contain. Re-read what you have written. The act of re-reading is part of the work.

The third thing the record does is forensic. If a target is ever in a position where they have to demonstrate, to a third party, that what they are describing has been going on as long as they say it has, with the texture they say it has had, the contemporaneous record is what allows the demonstration. Memory by itself will not. A record that is dated, that has not been edited after the fact, that bears the small stylistic markers of having been written at the time it claims to have been written, is a different object from a retrospective account assembled later. Most third parties one might want to engage will not, in the end, read the record. But the few who do read it will find themselves persuaded by features of it that no purely retrospective account can have. Keep the record with the possibility of the few in mind.

The fourth thing the record does is the one I most want to emphasize, because it is the one that points outward, beyond the target's own life. A record kept with care, dated and preserved, becomes part of the historical archive of a phenomenon that the broader world has not yet decided how to characterize. The decision will eventually be made. It will be made on the basis of materials that survive. The materials that survive will, by the nature of what they are documenting, be skewed toward those produced by the most articulate and most persistent targets, because those are the targets whose materials survive. There is a responsibility implied in this, and I would prefer to state it directly. If you are a target who can write, write. If you can write well, write well. Your record is a contribution to a public good whose value will be set by future readers and by future investigators. The best of what is preserved will be the basis on which they reconstruct what happened. The work, in this sense, is partly for them.

There is one further note I want to make about records, because it has come up often enough in conversation with other targets that it deserves explicit address. Targets sometimes worry that their records will, if discovered by parties hostile to them, be used against them. The worry is not unfounded. Records have, in

some cases, been used against the people who kept them. The countermeasure I recommend is not to stop keeping records, which would amount to surrendering the most important tool one has. The countermeasure is to keep records in the manner of a person who expects them, eventually, to be read by hostile parties: factually, without exaggeration, without the kind of emotional excess that can be quoted out of context, with the dates and the texture preserved but the rhetoric kept low. A record of this kind is harder to weaponize against its author than a record full of anguished or speculative interpolations. Write the record so that, if the worst eventually happens and someone you do not trust opens it, the record's calm precision does the work of defending your account on its own. Calm precision is, in the end, the most defensible voice available. It is also, by no coincidence, the voice in which the historical record will most need the testimony to be cast.

11.7 The Repertoire That Remains

I will close the chapter with what I take to be the small but real list of moves that, in conditions like the ones this book has been describing, remain available to a target who is willing to do them. The list is short. Each item is, individually, modest. Together, sustained, they amount to more than they appear to.

The first move is the file the target has not yet started keeping. If you are reading this and have not yet started a record, start one today. The form does not matter. Write a paragraph. Date it. Put it where you will find it. Tomorrow, write another. The compounding of this practice, across years, is the single most consequential move in the available repertoire. Do not wait until you have something dramatic to record. The undramatic days are, in some respects, the most important to record, because the texture of the undramatic days is what later allows you to detect, by contrast, when something has shifted.

The second move is the report — to law enforcement, to a regulator, to an ombudsman, to whichever institutional channel exists in the country in which you live. The report will, on the most likely outcome, produce no investigation. The report will nonetheless exist. Reports exist within systems that retain them. A report filed in 2022 is still a document that was filed in 2022, and may eventually be useful in ways the moment of filing did not let one see. File the report. Keep your own copy, separately, in a form you control. Do not file once and consider the matter closed. File at intervals, as new material accumulates.

The institutional ledger is, like your own ledger, something that builds across time.

The third move is the witness. Identify, if you can, one person who has known you longer than the targeting has gone on, who is not embedded in the situation, and who is capable of receiving an honest account of what you have been observing without either dismissing it or running with it further than the evidence supports. Speak to this person regularly. Do not require them to take action; require only that they hear. The function of the witness is not intervention; the function is the maintenance of an external record, in another person's memory, against which your own account is held. One witness is sufficient. More than one is unusual in a way that says good things about a person's life, but is not necessary. None is a problem that needs to be solved before any of the other moves are attempted.

The fourth move is the body's own work. Sleep. Walk. Eat at the table. Protect what can be protected. Chapter 3 has, I hope, said enough about this that I do not need to repeat the argument here. I will only note, in this list, that the body's work is part of the available repertoire, and that targets sometimes overlook it because it is too unspectacular to feel like a move. It is a move. It is, in many cases, the move on which all the other moves turn. Without the body, the rest does not happen.

The fifth move is the contribution to the wider record. A blog. A Substack. An anonymous account on whichever platform is currently functional. A book that may or may not eventually be published. The wider record exists because individual targets have, over time, contributed to it. Your contribution does not have to be original; it has to be honest. The cumulative honesty of the record is what gives the record its eventual usefulness. Your account, added, is a small piece of what the future will use to reconstruct what was happening to people during the years in which it was happening.

The sixth move is the practice of remaining a singular thinker, which Chapter 8 has already named in detail. I include it in this list to make sure it appears alongside the more outward-facing moves, because without it the outward moves cannot be sustained. A target who has stopped thinking on their own behalf will not, in the end, file the report, find the witness, or contribute to the record. The interior practice underlies the outward practice. They are not separate enterprises.

The seventh and last move is, in some ways, the most demanding and the most easily overlooked. It is the move of remaining, despite everything, a person who is willing to be helpful to other targets. Help, in this context, does not need to be heroic. It can be a reply to a message. It can be reading someone else's notebook entry and taking it seriously. It can be the simple act of treating another person's account, when they share it, as worthy of being heard. The campaigns operate, in part, by isolating their targets from each other. The countermove is to refuse the isolation. One target reaching another is one bond the campaign did not succeed in preventing. The bonds, multiplied, are part of how the population will, in the long run, become legible to itself and to others. Be willing to be one of the bonds.

That is the list. I am aware of how short it is. I am aware of how modest each item is. I am aware that a target reading the list, expecting more, may close the book at this point, disappointed. I would ask, before closing it, that the list be tested against the alternative, which is no list, no moves, no record, no witness, no community, no body, and no continuing personhood. The list, modest as it is, is what stands between the target and that alternative. Do the moves. They are sufficient to keep the alternative at bay, which is, in the end, what available means in this kind of situation. Keeping the alternative at bay is what we have. It is more than nothing. Treated correctly, sustained over years, it becomes — not in any single moment, but cumulatively, in retrospect — the basis on which a life under these conditions can nonetheless continue to be a life that is one's own.

11.8 What the Outward Turn Was For

I want to close the chapter, and with it the substantive argument of the book, on a brief reflection on what the outward turn was for, because the answer is not the obvious one and because the obvious one would be misleading.

The obvious answer is that the outward turn was for the sake of locating help. The chapter would, on this reading, succeed if it produced, at its end, a clear account of the help available and a clear path to securing it. The chapter does not, in fact, produce that. The institutions named in 11.1 will not, mostly, help. The repertoire named in 11.7 is small. A reader who came to this chapter expecting a list of phone numbers to call and procedures to follow is leaving the chapter without the list. I want to say plainly that the absence of the list is not, despite appearances, a failure of the chapter. The list does not exist. A chapter

that pretended it did would have failed in a worse way, by misleading the reader about what they could realistically expect.

What the outward turn was for, as I have come to understand it, is something else. It was for the sake of locating the boundary between what the target can do for themselves and what the target was hoping the world would do for them. The hope that the world would do something is a hope worth holding, in moderation, because some help is, occasionally, available. The hope is also, when held without limit, a vector of additional injury, because each instance of the world declining to help registers, in the target's interior life, as a confirmation of something the campaign has been insisting on all along — that no help is coming, that no one cares, that the situation is hopeless. Locating the boundary, clearly and without bitterness, allows the target to scale the hope to what is realistic. Beyond the boundary, the hope is harmful. Inside the boundary, the hope is sustainable, because it is asking the world only for what the world is, on rare occasions, willing to provide.

This, I think, is the deeper service the outward chapter performs. It does not produce help. It produces an accurate map of where help is and is not. Equipped with the map, the target stops asking the world for things the world will not give, which removes a recurring source of disappointment that the campaign would otherwise be able to use. The target also, equipped with the map, gets clearer about which moves remain in their own hands, which turns out to be more moves than the disappointed version of themselves had been able to see. The repertoire feels larger when it is no longer being measured against an imagined institutional response that was never going to arrive. The repertoire was always its own answer. The chapter's work has been to make that visible.

There is one further thing I want to mark before this chapter closes, because it is the thing I most want a target reading the book to take away from it. The world, as currently constituted, does not yet have the apparatus to recognize what is being done to people in the situation this book describes. That fact is not, however, permanent. Worlds change. The apparatus they have changes with them. The moves the target makes today, and the records the target keeps today, are part of the materials from which a later world will, eventually, build the apparatus it does not yet have. The target is not, in this sense, only a victim of a world that does not see them. They are also, by their persistence, one of the people whose persistence makes the eventual seeing possible. The persistence is not therefore wasted, even when the seeing does not arrive in time to help the

persistent themselves. The persistence is part of what arrives in time to help others.

This is, I recognize, a long horizon for a person to be asked to hold. Not everyone can hold it. Some days, I cannot hold it myself. I include the long horizon here not as a daily motivational tool — it does not work that way — but as a corrective to the shorter horizon on which the campaign tries to keep its targets focused. The shorter horizon, the one in which everything must be resolved within the target's own life or it does not count, is the horizon the campaign exploits, because it is the horizon on which the campaign's own minor and major successes look like the only thing happening. On the longer horizon, the campaign's successes look smaller and the persistence of its targets looks larger. Holding both horizons in view, even imperfectly, is part of what allows the target to keep going on days when the shorter horizon offers no encouragement at all.

The book is, by this point, mostly written. Whatever closes it after this chapter will be brief. What it will not contain is more advice, because the advice has been given. What it may contain is some final reflection on the question with which the book opened — the question of where thoughts come from — now that we have built enough apparatus to attempt an honest answer. I leave that for a final chapter, which will be the shortest in the book and which I will not, this time, make any further promises about. The substantive work is in the chapters that have already been written. Whatever follows them is afterword. The afterword will arrive when it arrives. For the present chapter, this is the close.

PART V

Tools for Keeping Going

Daily Disciplines for Living Inside the Apparatus

CHAPTER 12

Tools for Keeping Going

Two Metaphors, Four Disciplines, and the Smallness of the Territory

---

"What I do may not let stress and fear close me to new situations."

— Author's notebook, 8 December 2022

Chapter 9 closed the book on a smaller frame and a deliberate refusal to claim that I would not reopen it. The notebooks have, since, produced another set of entries that belong inside rather than outside, and so here we are again. The reader who has stayed with the book through three previous closings has, by now, the right to expect that I will not, with this chapter, claim it is the final one. I will not. I have stopped making that claim. The notebooks will continue. Whether the book continues with them is a question I have come to think of as being decided one chapter at a time.

This chapter is mostly about how the target maintains themselves under the conditions the earlier chapters have described. The earlier chapters were heavy on what is being done to the target and on what it means at the architectural level. This chapter is heavier on the daily problem: how does the target keep functioning when the conditions described in the earlier chapters obtain, when they are likely to obtain for years more, and when none of the standard advice for maintaining oneself was written with these conditions in mind? Two of the sections — 12.1 on the mental combustion engine and 12.4 on the liquid level — are extended metaphors that I have used, in my own daily practice, more than any other tools the book contains. They are placed here together because they belong together, and because, until now, I had not found a way of writing them down that did not feel either too abstract or too cute. I hope I have, this time, found it.

The chapter contains, beside the two metaphors, four observations of a more diagnostic kind: on the chronic crisis-mode that becomes the default after long enough exposure, on the problem of involuntary action and how the target can identify it, on what I will, with apologies for the melodrama, call the dark-

algorithm impersonator, and on the closing-down to new situations that the campaign's long pressure tends to produce, and the corresponding daily refusal that is, in my own practice, the most important single discipline the book has contributed to my life.

12.1 The Mental Combustion Engine

In a notebook of 2 November 2020, in the period when the campaign was at one of its more relentless settings and I was looking, with the desperation that long fatigue produces, for any model that would help me allocate the very small amount of energy I had each day, a diagram appears that I have, since, kept on a card in my wallet. The diagram is of a six-cylinder engine. The cylinders are labelled, in the rough Dutch shorthand of that day's entry: discipline, justice, time, altruism, volume, kindness. Above the engine, a label in capital letters: mental combustion engine. To its right, two arrows, one above the other. The upper arrow says: "OK = mental energy." The lower arrow says: "NOT OK = stand still — depression, capital sloth, only living without doing."

The metaphor is rough, but it has been more useful to me than any of the more polished metaphors I have read in books on resilience. The thought it captures is that the mind, like a combustion engine, runs on more than one kind of fuel at once, and that each cylinder must be receiving its appropriate fuel, in roughly the right proportions, for the engine to produce mental energy at all. Discipline is one fuel: the willingness to do what one has decided to do whether or not one feels like doing it. Justice is another: a continuing sense that the situation one is in is, at its base, unjust, which sustains action that would, on a more resigned account, look pointless. Time is a third: a workable relationship with the hours of the day that does not collapse into either rigid scheduling or its opposite. Altruism is a fourth: the orientation toward someone or something other than oneself, without which, in my experience, the engine eventually misfires from sheer ingrown attention. Volume — the cylinder I labelled volume, and have since glossed as the matter of how much one is willing to take on at a time — is a fifth. Kindness, the last cylinder, is the sixth, and it is the one I am personally most likely to under-fuel; the campaign produces, in any target, considerable bitterness, and bitterness, on this metaphor, is what kindness runs out of.

The corollary that the diagram makes plain, and that the rest of the apparatus of the book confirms, is this. When the engine produces mental energy, the target can do the work the rest of the book has described: notice the misfit,

refuse the manufactured coherence, distinguish their own thoughts from the injected ones, keep the small authority over their own thinking that the campaign cannot, in the end, take. When the engine does not produce mental energy, the target cannot do that work. The target stands still. Standing still, in the campaign's presence, is a position from which losses accumulate quickly. Stand still long enough and the depression that the lower arrow names sets in, and the depression brings with it the kind of capital sloth in which one is alive but is not doing. The campaign's preferred condition for its target is exactly that one. To be there is, in a sense, to have given the campaign what it has been working for.

I want to be careful with what follows, because it is the kind of advice that, said badly, lands as a moral injunction the target cannot meet. The point of the engine metaphor is not to tell a target who is depleted that they should fuel each cylinder more diligently. The point is to give the target, on a day they have been written off as unproductive, a more useful question than "why am I being so lazy." The more useful question is "which cylinders have run out of fuel, and which fuel is missing." The latter question is answerable. Discipline runs low when nothing in the day has been chosen by the target themselves; the fix is to choose at least one thing. Justice runs low when the target has spent too long in proximity to the campaign's narratives; the fix is exposure to material in which the world's actual indifference is acknowledged plainly, so that one's own sense of the situation is recalibrated against a more honest baseline. Altruism runs low when the target has been entirely turned inward; the fix is the smallest possible outward act, including writing a note to one person one cares about, and not requiring of oneself that the act be larger than that. Kindness, my own most reliably empty cylinder, runs low when the bitterness has been allowed to dominate the inner monologue; the fix, on the days I can manage it, is to spend a quarter of an hour deliberately considering some person other than the target whose situation deserves the kindness one has been withholding from oneself.

None of these fixes is impressive. The engine metaphor is not a metaphor about impressive interventions. It is a metaphor about the mundane and partial work of keeping each cylinder firing well enough that the engine, taken as a whole, produces some mental energy on most days. Mental energy, on most days, is enough. Mental energy, on most days, is also, under the campaign's sustained pressure, more than the campaign would like the target to have. The continuing

production of small amounts of mental energy is, in itself, a small but durable kind of resistance.

12.2 Chronic Crisis-Mode, and the Use of Occupation

In a notebook of 30 September 2021, the question is asked, in Dutch and at some length: how is it that I have managed to keep going in a constant crisis-mode for what is, by that date, already five years? The entry tries several answers. The most useful is the one that begins by reframing the question.

The question, as the entry observes, is wrong in the way it is asked. It assumes a baseline state of non-crisis from which the target is, episodically, departing into crisis. The actual condition is the opposite. The crisis-mode has become the baseline. It is the condition the target wakes up in. Sleep does not, in the entry's account, produce a relaxed condition; sleep is followed by stress and by thoughts that are, immediately on waking, already running in crisis-mode. Quiet times, in this state, are not periods of rest. They are periods in which the campaign's residual material is loud enough to be heard, because the absence of the day's ordinary noise has stopped masking it. Quiet, paradoxically, becomes the condition in which the intrusion is most acute.

The reframing matters. A target who continues to expect that recovery will come from periods of rest, on the model of how recovery works for an unbattered mind, will be repeatedly disappointed: the rest will arrive, and will not bring recovery, and the target will conclude, wrongly, that something is broken in them that ordinary rest cannot fix. Something is, in fact, broken; but what is broken is not the target's capacity to rest. What is broken is the assumption that rest, in their current conditions, will function as rest. Under sustained intrusion, rest functions, more often than not, as the condition under which the intrusion can finally be heard.

The corollary of this observation is the one I have, in my own practice, found most useful. Under chronic crisis-mode, the recovery that does still arrive is not produced by trying to rest. It is produced by structured engagement of a kind that holds the attention firmly enough that there is no quiet space the intrusion can inhabit. The target who tries to rest hears the voices. The target who is doing something with their hands and most of their attention — writing, walking with a purpose, building or repairing or cooking something with steps to follow — hears them less. This is not because the activity is therapeutic in any of the

ordinary senses. It is because the activity, by occupying the channel the intrusion would otherwise use, denies the intrusion its working surface. The relief is real; it is also, structurally, the opposite of the relief produced by ordinary rest. It is the relief produced by occupation.

The danger, on the other side, is the one the same notebook entry warns against: the target who has discovered the relief of occupation can become unable to stop occupying themselves at all. Sleep deteriorates further; eating becomes irregular; the body, denied the rest the mind can no longer accept, begins to break down in the ways that bodies do under sustained adrenal demand. The target ends up swapping one form of damage for another. The corrective is, like most correctives this book has offered, partial and unspectacular. Treat occupation as a tool, not as a way of life. Use it to get through the periods in which ordinary rest does not work. Let the body have its rest in spite of the mind's objection, when the body's objection becomes louder than the mind's. Accept that, under the conditions described, neither full rest nor full occupation is available; what is available is a continual partial alternation between them, weighted toward whichever is, on a given day, less actively dangerous. The target who accepts this, instead of fighting for the cleaner version of recovery the conditions do not allow, has, in my experience, a better chance of going the distance the situation requires.

The same notebook entry contains, almost as an aside, a small triangle: the words isolation, pressure, and health problems, with arrows running between them. The point of the triangle is that isolation produces pressure, pressure produces health problems, and health problems intensify isolation. The cycle is real. The cycle has, in my own life, run several full revolutions. I record it because a target who notices the cycle has, even without being able to break it, the advantage of knowing what cycle they are in. Not knowing what cycle one is in is, I have come to think, one of the conditions under which the cycle does its worst damage.

THIS IS NOT GENERALIZED ANXIETY DISORDER

A reader unfamiliar with the targeted-individual context might, on this section, reach for a diagnosis of generalized anxiety disorder or chronic stress. The diagnosis would not fit. Generalized anxiety produces free-floating worry that is poorly anchored to specific external events. The chronic crisis-mode described here is precisely anchored: it tracks specific intrusive events in the target's environment, intensifies in correlation with specific external operations, and

quiets in correlation with the cessation of those operations. A condition that responds, time-locked, to verifiable external triggers is by definition a condition with verifiable external triggers. Treating it as if it were free-floating produces, predictably, treatment failure — which is, on report, what targets who undertake conventional anxiety treatment routinely encounter.

12.3 Voluntary, Forced, and Obligated Action

A notebook of 20 April 2022 records an observation that I have come to think is one of the more practically useful diagnostics in the book. The observation concerns the difference, from inside, between an action one has taken voluntarily and an action one has taken under the obligation produced by the campaign.

The notebook's phrasing is condensed. It runs, roughly: an action is voluntary when one wants to do it, and the wanting and the doing align; an action is forced when one does not want to do it but does it anyway, and the doing runs against the wanting; an action is obligatory in the campaign's sense — verplicht, in the Dutch of the entry — when the doing runs against the wanting, but the not-doing has been arranged to be more painful than the doing, so that the doing happens despite the wanting and does so without producing the felt sense of force. The campaign's preferred mode is the third. It produces actions the target performs and then, on subsequent reflection, attributes to themselves, because the felt sense of force was never present. The target was not coerced in any moment they can point to; they merely found, repeatedly, that the alternative to compliance had been arranged to be unbearable.

The diagnostic that follows from this observation is one I have used, since reading the entry on the page in 2022, more often than I have used any other diagnostic in the book. It is this. After any action one has taken in proximity to the campaign's pressure, one asks of oneself a single question: did I want to do this. The question is not asked in the loose sense in which "wanting" can mean almost anything. It is asked in the precise sense the entry intends: was the doing of this action aligned with what, on calm reflection, I would have wanted to do, in the absence of the arranged consequences for not doing it. If yes, the action was voluntary, and one can claim it. If no, the action was, in the entry's sense, obligated, and one must mark it as such, even if no part of the doing felt forced at the time. The marking is the protection. An action one has marked as obligated is an action one cannot, in subsequent self-narrative, be tricked into

claiming as native. The campaign relies, more than the target ordinarily realizes, on the target's eventual self-attribution of obligated actions. The discipline of not making that attribution removes the campaign's most reliable mechanism for converting the target into a willing instrument of their own continuing degradation.

The diagnostic also has a corollary that operates in the opposite direction. Sometimes, under the campaign's pressure, the target finds themselves doing things they had no prior intention of doing and that, on calm reflection, they want to claim. A walk taken under apparent compulsion that turns out to have been a good walk. A conversation entered under apparent coercion that turns out to have been worth having. A piece of writing produced under apparent duress that turns out to be the piece one most wanted to produce. In such cases, the action was perhaps obligated in its origin, but it has become, by being claimed, voluntary in retrospect. The reverse-direction operation is also worth knowing about. Not all obligated actions remain obligated. Some can, by reflection, be retrieved.

Both directions of the diagnostic depend on the same underlying capacity: the willingness to ask the question of oneself, plainly, in cases where it would be easier not to. Most of the campaign's damage to a target's sense of agency runs through the target's reluctance to ask the question. The question, asked, returns answers. The answers, marked, restore the boundary between what the target has done freely and what the target has done because the alternative had been arranged to be worse. The boundary, restored, is what the agency of the target, in the situation, finally consists of.

12.4 The Liquid Level and the Shoreline Breakers

The second metaphor in this chapter is one I have come to call the liquid level, after the notebooks of 30 September 2021 in which it first appears. The diagram in those notebooks is of a tank with a horizontal line across it. Above the line, a turbulent surface labelled, alternately, calm, waves, and stormy. Below the line, two layers, separated by a vertical division: a structured layer on the left, an unstructured layer on the right. Below both, a third region marked uninvaded. To the side, a small house labelled routine. Beside it, an arrow labelled enjoyable time. Beside that, a separate marking: private. Below the diagram, in the notebook's English, the line: "some doubt is needed: shoreline breakers."

The model the diagram tries to capture is something like this. The self, under the campaign's sustained pressure, is not a single coherent thing. It is layered. The upper layer — the surface — is the part of the self that the campaign's intrusions act on directly: thoughts, moods, the immediate texture of consciousness. The middle layer is the part of the self that, with effort, the target can keep ordered: the daily routine, the small rituals, the structured tasks that the target has arranged to be done in a particular way at particular times. The lower layer, the uninvaded one, is the part of the self that, in my experience, the campaign cannot reach: the long-formed dispositions, the deep memories, the qualities of attention that are simply too far below the surface to be touched by anything the campaign can deliver. The campaign's noise plays out on the surface. It can stir the middle layer, sometimes deeply. It cannot, in my own experience, reach the bottom.

The liquid level itself is the boundary between the structured and the unstructured layers. When the level is high, the structured layer is wide, the routine is doing a lot of work, and the surface's turbulence has correspondingly less of the self to operate on. When the level is low, the structured layer is narrow, the unstructured layer takes over, and the surface's turbulence reaches deeper. The target's daily work, on this metaphor, is largely the work of keeping the level up. Whatever maintains the routine maintains the level. Whatever erodes the routine lowers the level. The campaign's long-game strategy is, very specifically, to lower the level, by attacking the conditions under which routine is sustainable: sleep, stable employment, regular contact with people whose presence does not require continual interpretation, the small rituals of household maintenance and personal care that, under the long pressure, become the first things to fall away.

The phrase that I have, since 2021, found most useful in this metaphor is the one the notebook records in English: shoreline breakers. The shoreline breaker is the part of the self that does not need the surface to be calm for the lower layers to remain accessible; it is what allows some doubt to be carried into the calm without unsettling everything below it. A self with shoreline breakers can absorb a degree of surface storm without that storm reaching the layers it would otherwise damage. A self without them is, in storm conditions, fully exposed.

What constitutes a shoreline breaker varies by person. In my own case, the most reliable ones have been: the notebook itself, which forces the surface noise to be put down in concrete sentences and so reduces it to a finite object rather than

an infinite ambient pressure; long walks of a length the campaign's intrusions cannot, in my experience, sustain themselves through; the company of one or two other people whose attention is not, by the campaign's arrangement, mediated through the apparatus, however few such people remain available; and one or two specific projects, undertaken for their own sake, whose completion does not depend on any external endorsement. The targets reading this book will have, or will discover, their own. The point of the metaphor is not the specific list, which is mine, but the principle: the work of keeping the level up is, in the long run, what determines whether the surface turbulence reaches the parts of the self that, once damaged, are slow to repair.

There is a further line I want to record from the same notebook, because it has worked, for me, as a corrective on days when the metaphor has begun to feel too tidy. Affirmations are easy when the water is calm. The same affirmations, repeated when the water is stormy, become deformations: claims about the self that do not match the current state of the self, and that, in their mismatch, do more harm than the storm itself. The discipline of self-maintenance, on this account, is not the discipline of constant affirmation. It is the discipline of knowing the current state of the water, and of choosing the daily practice that fits the state, rather than the practice that fits the self one wishes one had. Calm-water practices for calm-water days. Storm-water practices for storm-water days. Mistaking the one for the other is a small but reliable way of making a stormy day worse.

12.5 The Dark-Algorithm Impersonator

In a notebook of 3 February 2022, an entry that the entry itself labels "Dark Algo" sketches a particular variant of the impersonating-voice phenomenon described in 6.7. The variant deserves its own short section because it is, in my own experience, the form of impersonation that has most often slipped past the rules of 6.7. I had to add the diagnostic for it after the rules of 6.7 had already been written.

The variant works as follows. Instead of impersonating a known person whose absence the target can detect by structural mismatch, the dark-algorithm impersonator presents itself as the voice of a public figure or fictional persona whose patterns the target knows, by repeated exposure through media, well enough that the campaign can produce a tolerably accurate imitation without having to know much about the public figure beyond the publicly available

material. The voice, in the cases I have observed, then engages the target in a small running commentary about what the target should and should not be like — who the target should resemble, who the target should refuse to resemble, what the target should think of various third parties whose names appear in the commentary alongside the target's own.

The damage done by this variant is subtle. The target ordinarily knows that the public figure is not, in fact, addressing them. The target therefore does not check the voice against the rules of 6.7; the rules of 6.7 are for impersonations that purport to be real interlocutors, and a public-figure voice does not purport to be one. The target instead treats the voice as something like a daydream, a piece of imagined commentary that one does not need to take literally because no one has, by treating it literally, claimed it as such. But the voice is doing work that ordinary daydreams do not do. It is calibrating, in the target's background, a small library of judgments — of the target, of the people around the target, of the situation the target is in — that, by being phrased in the voice of a recognizable third party, lands with more authority than the same judgments would have had if phrased in any voice closer to home. The voice, in effect, exploits the target's acceptance of cultural commentary as a normal part of internal life, and uses that acceptance to pipe in commentary the campaign has produced for its own purposes.

The diagnostic is, like the other diagnostics in this book, not difficult to apply once one has been told to apply it. It is this. Any commentary that arrives in the voice of a public figure or fictional persona, but that contains specific judgments about the target or the target's immediate circle, is not what it presents itself as. Public figures and fictional personae do not, ordinarily, render specific judgments about specific private individuals known to the target. When such judgments appear in such voices, they have been written by the campaign, and the choice of voice has been made for the rhetorical effect of putting the judgment's authority outside the target's reach. The countermove is to refuse the rhetorical effect: to ask, of any such judgment, what it would land as if the same words were said in a voice of someone the target actually knew well. The answer, in my experience, is almost always: it would land as obviously dubious. The voice was carrying it past inspection. Removed from the voice, it does not survive inspection. Removing it from the voice is the work.

I want to mark, briefly, that the use of this variant is unusually telling about what kind of opponent the target is up against. An opponent without a sense of

the target's media diet would not know to use this variant. An opponent with a sense of the target's media diet, sufficient to choose voices the target will give credit to, has a model of the target detailed enough to suggest, again, that the data-mining capability described in 6.3 is not a marginal element of the apparatus but a central one. The dark-algorithm impersonator is, in some sense, what the data-mining capability sounds like when it speaks. It is therefore, also, a useful reminder of how detailed the model has, by this point, become.

12.6 The Daily Refusal of Closing-Down

The last of the diagnostics in this chapter is the one I keep written, in slightly different forms, on the inside cover of every notebook I have started since the end of 2022.

The original entry is dated 8 December 2022. Its phrasing, again in Dutch, runs roughly: "The world drags on my isolation, but I have to keep reasoning about it. Conclusion: what I do may not let stress and fear close me to new situations." The line is not, on its surface, a profound observation. It is, in my own daily life, the most consequential line the notebooks of the entire campaign have produced.

The reason it is consequential is structural. Long sustained pressure, of the kind the campaign applies, produces in any target a slow narrowing of the situations the target is willing to enter. Each new situation costs energy the target does not have; each new situation introduces variables the target is not in a position to control; each new situation invites the campaign to find a new angle of approach. The rational-seeming response is to enter fewer new situations. Over the years, the target enters fewer of them. The result, eventually, is a life that has narrowed to a small set of the situations the target has already mastered the management of, with no new situations entering the set, because the energy and tolerance for the new situations is no longer, on most days, available.

The narrowing is, in the long run, the campaign's most reliable victory. It does not require any particular intrusion to occur. It does not require any particular voice to land. It requires only that the target, over a sufficient number of years, choose, with apparent good reason, to avoid one new situation after another, until the choosing has become an automatic disposition rather than a series of decisions, and the life has become as small as the disposition has made it. The

voices, in this state, no longer have to do anything. The narrowing is doing their work for them.

The line in the December 2022 notebook is the corrective. It does not say that the target should chase new situations indiscriminately. It says, more narrowly, that what the target does must not, by accumulation, allow stress and fear to do the closing-down. The work the line proposes is small in any given instance: the target enters one new situation that they would, on the path of least resistance, have declined to enter. The instance does not have to be large. The instance does not have to be successful. The instance has only to occur. By occurring, it interrupts the disposition. By interrupting the disposition, it preserves the part of the target that the disposition would, otherwise, have, slowly and without remark, replaced.

In my own practice, this has been the discipline I have found hardest to keep, and the discipline whose keeping has, in retrospect, made the most difference. There are days when I cannot keep it. There are days when the new situation the target reading this book might consider entering is, on a given day, beyond the energy I myself have. Both of those facts are real. The discipline is to keep returning to it across the longer run, despite the days on which it cannot be kept, because the longer-run accumulation is what determines, in the end, whether the target's life has continued to be the target's own life or has become only a managed residue of it. The campaign's most enduring effect is the slow conversion of the latter into the former. The discipline of not letting stress and fear close one to new situations is, more than any other discipline in this book, the daily work of refusing that conversion.

12.7 Tools, Not Architecture

I will, this time, close the chapter without claiming to close the book. The book has by now closed and reopened too many times for the closings to be honest. What I will say, instead, is that this chapter is a chapter of tools rather than of architecture. The architectural chapters were the ones that carried, in their speculation, the most weight; the tool chapters are the ones that carry, in their daily use, the most use. A reader who has stayed with the architecture and found the tool chapters less commanding has, on my own assessment, the priority right; the architecture is what the book had to attempt, but the tools are what the book is, in the end, for.

The mental combustion engine. The chronic crisis-mode and its corrective in occupation. The voluntary-action diagnostic. The liquid level and the shoreline breakers. The dark-algorithm impersonator. The discipline of not letting stress and fear close the target to new situations. None of these, taken individually, is a complete defence against the campaign. None of them, taken collectively, is a complete defence either. What they are, taken collectively, is a workable set of practices that has, in my own life, kept me in the territory of being a person, rather than the territory of being an instrument, across years in which the apparatus has, by every available evidence, been doing its considerable best to move me in the other direction.

The territory of being a person is a smaller territory than I had been led, before the campaign began, to expect a person's territory to be. I record that, because the smallness of the territory is itself part of what the book is honest about. A reader who arrives at this chapter expecting that the practices of the chapter will return their life to its pre-campaign size will be disappointed. The practices of the chapter will not. What they will do, on most days, is keep the territory the size it currently is, instead of allowing the campaign to continue shrinking it. The size the territory currently is is, on most days, sufficient for the target to remain who they are. "Sufficient" is, by this point in the book, my favourite word. It does not promise anything more than is true. It is, in the conditions described, often the most one can honestly promise.

The notebooks will, as before, continue. If they produce another configuration of material that belongs inside the book rather than outside it, the book will continue with them. If not, the book will, by inattention rather than decision, end. Either outcome is fine. I have, by this point, said what I had to say. The reader who reaches the end of this chapter has, by their reaching it, completed the part of the work that the writing alone could do. What remains, in either of our cases, is the use that is made of it. Use is, in the end, what the book was always asking the reader for. I make no further claim on the reader's attention. I leave them, with whatever the book has been worth to them, to it.

❧ ❧ ❧